



US009178558B2

(12) **United States Patent**
Lea et al.

(10) **Patent No.:** **US 9,178,558 B2**
(45) **Date of Patent:** ***Nov. 3, 2015**

(54) **BACKHAUL RADIO WITH HORIZONTALLY OR VERTICALLY ARRANGED RECEIVE ANTENNA ARRAYS**

25/005 (2013.01); **H04B 1/50** (2013.01); **H04B 7/0408** (2013.01); **H04B 7/0695** (2013.01);
(Continued)

(71) Applicant: **CBF Networks, Inc.**, San Jose, CA (US)

(58) **Field of Classification Search**

(72) Inventors: **David Andrew G. Lea**, Vancouver (CA);
Kevin J. Negus, Jackson, WY (US)

None

See application file for complete search history.

(73) Assignee: **CBF Networks, Inc.**, San Jose, CA (US)

(56)

References Cited

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

U.S. PATENT DOCUMENTS

This patent is subject to a terminal disclaimer.

D307,274 S 4/1990 Sasaki et al.
D335,128 S 4/1993 Soren

(Continued)

(21) Appl. No.: **14/632,624**

FOREIGN PATENT DOCUMENTS

(22) Filed: **Feb. 26, 2015**

EP 2745539 6/2014
EP 2767102 8/2014

(65) **Prior Publication Data**

(Continued)

US 2015/0194999 A1 Jul. 9, 2015

OTHER PUBLICATIONS

Related U.S. Application Data

“AccessGate—RAN Optimization for Mobile Backhaul Systems,” Product Data Sheet, Memotec, 2009, Montreal, Quebec, Canada, 2 pages.

(Continued)

(63) Continuation of application No. 14/336,958, filed on Jul. 21, 2014, now Pat. No. 9,001,809, which is a continuation of application No. 13/898,429, filed on May 20, 2013, now Pat. No. 8,824,442, which is a continuation of application No. 13/536,927, filed on Jun. 28, 2012, now Pat. No. 8,467,363, which is a continuation-in-part of application No. 13/371,366, filed on Feb. 10, 2012, now Pat. No. 8,311,023, which is a continuation of application No. 13/212,036, filed on Aug. 17, 2011, now Pat. No. 8,238,318.

Primary Examiner — Hassan Phillips

Assistant Examiner — Gautam Sharma

(74) *Attorney, Agent, or Firm* — Nixon Peabody LLP

(51) **Int. Cl.**

H04W 4/00 (2009.01)

H04B 1/401 (2015.01)

(Continued)

(57)

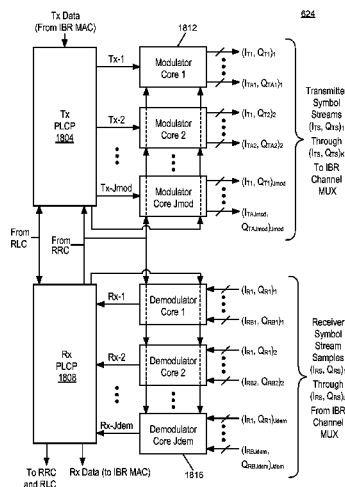
ABSTRACT

A intelligent backhaul radio have an advanced antenna system for use in PTP or PMP topologies. The antenna system provides a significant diversity benefit. Antenna configurations are disclosed that provide for increased transmitter to receiver isolation, adaptive polarization and MIMO transmission equalization. Adaptive optimization of transmission parameters based upon side information provided in the form of metric feedback from a far end receiver utilizing the antenna system is also disclosed.

69 Claims, 56 Drawing Sheets

(52) **U.S. Cl.**

CPC **H04B 1/401** (2013.01); **H01Q 1/246** (2013.01); **H01Q 21/24** (2013.01); **H01Q 21/29** (2013.01); **H01Q 25/00** (2013.01); **H01Q**



- (51) **Int. Cl.**
- | | | | | |
|-------------------|-----------|-----------------|---------|-------------------|
| <i>H01Q 21/24</i> | (2006.01) | 7,848,241 B2 | 12/2010 | Natarajan et al. |
| <i>H04W 88/08</i> | (2009.01) | 7,948,942 B2 | 5/2011 | Zhang et al. |
| <i>H04L 5/00</i> | (2006.01) | 7,978,144 B2 | 7/2011 | Tanabe et al. |
| <i>H04L 25/03</i> | (2006.01) | D643,025 S | 8/2011 | Podduturi |
| <i>H04B 10/25</i> | (2013.01) | 8,078,109 B1 | 12/2011 | Mulcay |
| <i>H01Q 1/24</i> | (2006.01) | 8,175,535 B2 | 5/2012 | Mu |
| <i>H01Q 21/29</i> | (2006.01) | 8,238,318 B1 | 8/2012 | Negus |
| <i>H01Q 25/00</i> | (2006.01) | 8,249,527 B2 | 8/2012 | Rybicki |
| <i>H04B 7/04</i> | (2006.01) | 8,300,590 B1 | 10/2012 | Negus et al. |
| <i>H04B 7/06</i> | (2006.01) | 8,311,023 B1 | 11/2012 | Negus |
| <i>H04B 7/08</i> | (2006.01) | 8,385,305 B1 | 2/2013 | Negus et al. |
| <i>H04B 7/10</i> | (2006.01) | 8,422,540 B1 | 4/2013 | Negus et al. |
| <i>H04B 1/50</i> | (2006.01) | 8,467,363 B2 | 6/2013 | Lea et al. |
| <i>H04L 27/01</i> | (2006.01) | 8,502,733 B1 | 8/2013 | Negus et al. |
| <i>H04L 27/26</i> | (2006.01) | 8,638,839 B2 | 1/2014 | Negus et al. |
| <i>H04L 1/00</i> | (2006.01) | D704,174 S | 5/2014 | Negus et al. |
| <i>H04L 1/18</i> | (2006.01) | 8,761,100 B2 | 6/2014 | Negus et al. |
| | | 8,811,365 B2 | 8/2014 | Negus |
| | | 8,824,442 B2 | 9/2014 | Lea et al. |
| | | 8,830,943 B2 | 9/2014 | Negus et al. |
| | | 8,872,715 B2 | 10/2014 | Lea et al. |
| | | 8,928,542 B2 | 1/2015 | Lea et al. |
| | | 8,942,216 B2 | 1/2015 | Negus et al. |
| | | 8,948,235 B2 | 2/2015 | Negus et al. |
| | | 9,001,809 B2 | 4/2015 | Lea et al. |
| | | 9,049,611 B2 | 6/2015 | Negus et al. |
| | | 9,055,463 B2 | 6/2015 | Negus et al. |
| | | 2001/0030939 A1 | 10/2001 | Vijayan et al. |
| | | 2001/0035844 A1 | 11/2001 | Reece et al. |
| | | 2001/0050927 A1 | 12/2001 | Johnson |
| | | 2002/0060993 A1 | 5/2002 | Dent |
| | | 2002/0064141 A1 | 5/2002 | Sakakura |
| | | 2004/0006573 A1 | 1/2004 | Takashi |
| | | 2004/0137924 A1 | 7/2004 | Herscovich et al. |
| | | 2004/0151238 A1 | 8/2004 | Masenten |
| | | 2004/0165650 A1 | 8/2004 | Miyazaki et al. |
| | | 2005/0026616 A1 | 2/2005 | Cavalli et al. |
| | | 2005/0075078 A1 | 4/2005 | Makinen et al. |
| | | 2005/0219140 A1 | 10/2005 | Browne et al. |
| | | 2006/0056365 A1 | 3/2006 | Das et al. |
| | | 2006/0079290 A1 | 4/2006 | Seto et al. |
| | | 2006/0141929 A1 | 6/2006 | Lockie et al. |
| | | 2006/0164271 A1 | 7/2006 | Hirt et al. |
| | | 2006/0209746 A1 | 9/2006 | Asai et al. |
| | | 2006/0253526 A1 | 11/2006 | Welch et al. |
| | | 2006/0273952 A1 | 12/2006 | Krikorian et al. |
| | | 2007/0097899 A1 | 5/2007 | Larsson et al. |
| | | 2007/0146225 A1 | 6/2007 | Boss et al. |
| | | 2007/0155431 A1 | 7/2007 | Munzner et al. |
| | | 2007/0243878 A1 | 10/2007 | Taira et al. |
| | | 2007/0264935 A1 | 11/2007 | Mohebbi |
| | | 2008/0014948 A1 | 1/2008 | Scheinert |
| | | 2008/0080364 A1 | 4/2008 | Barak et al. |
| | | 2008/0090575 A1 | 4/2008 | Barak et al. |
| | | 2008/0159212 A1 | 7/2008 | Zhang et al. |
| | | 2008/0181282 A1 | 7/2008 | Wala et al. |
| | | 2008/0240307 A1 | 10/2008 | Wang et al. |
| | | 2008/0242232 A1 | 10/2008 | Zavadsky et al. |
| | | 2008/0274745 A1 | 11/2008 | Barak et al. |
| | | 2009/0010238 A1 | 1/2009 | Barak et al. |
| | | 2009/0052411 A1 | 2/2009 | Leung et al. |
| | | 2009/0067526 A1 | 3/2009 | Ratermann et al. |
| | | 2009/0111463 A1 | 4/2009 | Simms et al. |
| | | 2009/0143017 A1 | 6/2009 | Barak et al. |
| | | 2009/0252134 A1 | 10/2009 | Schlicht et al. |
| | | 2009/0274130 A1 | 11/2009 | Boch |
| | | 2009/0304055 A1 | 12/2009 | Nino et al. |
| | | 2009/0312022 A1 | 12/2009 | Viorel et al. |
| | | 2009/0323621 A1 | 12/2009 | Touboul et al. |
| | | 2010/0009694 A1 | 1/2010 | Fischer |
| | | 2010/0033396 A1 | 2/2010 | Tanabe et al. |
| | | 2010/0046439 A1 | 2/2010 | Chen et al. |
| | | 2010/0056205 A1 | 3/2010 | Fuss |
| | | 2010/0157970 A1 | 6/2010 | Gotman et al. |
| | | 2010/0202391 A1 | 8/2010 | Palanki et al. |
| | | 2010/0261423 A1 | 10/2010 | Stanforth et al. |
| | | 2010/0272006 A1 | 10/2010 | Bertrand et al. |
| | | 2010/0309048 A1 | 12/2010 | Polisetty et al. |
| | | 2011/0044279 A1 | 2/2011 | Johansson et al. |
- (52) **U.S. Cl.**
- CPC *H04B 7/088* (2013.01); *H04B 7/10* (2013.01); *H04B 10/25* (2013.01); *H04L 5/0007* (2013.01); *H04L 25/03828* (2013.01); *H04L 27/01* (2013.01); *H04L 27/2636* (2013.01); *H04W 88/08* (2013.01); *H04L 1/0001* (2013.01); *H04L 1/004* (2013.01); *H04L 1/18* (2013.01); *H04L 27/2601* (2013.01)
- (56) **References Cited**
- U.S. PATENT DOCUMENTS
- | | | |
|--------------|---------|---------------------|
| D349,897 S | 8/1994 | Soren et al. |
| D373,771 S | 9/1996 | Messelhi |
| 5,579,367 A | 11/1996 | Raymond et al. |
| D376,367 S | 12/1996 | Mailandt |
| D376,600 S | 12/1996 | Vallilee et al. |
| 5,604,744 A | 2/1997 | Andersson et al. |
| 5,809,422 A | 9/1998 | Raleigh et al. |
| D401,592 S | 11/1998 | Nishimura et al. |
| 5,890,055 A | 3/1999 | Chu et al. |
| RE36,591 E | 2/2000 | Hayashi et al. |
| 6,253,060 B1 | 6/2001 | Komara et al. |
| 6,310,584 B1 | 10/2001 | Reece et al. |
| D455,420 S | 4/2002 | Arpe |
| 6,377,217 B1 | 4/2002 | Zhu et al. |
| 6,462,710 B1 | 10/2002 | Carson et al. |
| 6,463,303 B1 | 10/2002 | Zhao |
| 6,486,828 B1 | 11/2002 | Cahn et al. |
| 6,529,172 B2 | 3/2003 | Zimmerman |
| 6,912,195 B2 | 6/2005 | Vook et al. |
| D507,263 S | 7/2005 | Narita |
| 6,985,123 B2 | 1/2006 | Gottl |
| 6,995,712 B2 | 2/2006 | Boyanov |
| 7,027,837 B1 | 4/2006 | Uhlik et al. |
| 7,171,223 B2 | 1/2007 | Herscovich et al. |
| 7,221,722 B2 | 5/2007 | Thomas et al. |
| 7,260,141 B2 | 8/2007 | Bierly et al. |
| 7,280,082 B2 | 10/2007 | Theobald et al. |
| 7,292,663 B1 | 11/2007 | Van Wechel et al. |
| D565,534 S | 4/2008 | Ingalsbe et al. |
| 7,394,439 B1 | 7/2008 | Johnson et al. |
| 7,403,501 B2 | 7/2008 | Bordonaro et al. |
| D595,274 S | 6/2009 | Skottke et al. |
| D596,627 S | 7/2009 | Cho et al. |
| 7,587,177 B1 | 9/2009 | Kwong |
| 7,593,729 B2 | 9/2009 | Barak et al. |
| 7,599,290 B2 | 10/2009 | Dos Remedios et al. |
| D604,724 S | 11/2009 | Vorreiter et al. |
| 7,616,554 B2 | 11/2009 | Asai et al. |
| 7,620,370 B2 | 11/2009 | Barak et al. |
| D619,571 S | 7/2010 | Lee |
| 7,756,519 B2 | 7/2010 | Barak et al. |
| D622,256 S | 8/2010 | Lockenwitz |

(56)

References Cited**U.S. PATENT DOCUMENTS**

2011/0070855	A1	3/2011	Mariotti
2011/0085525	A1	4/2011	Patini
2011/0103292	A1	5/2011	Pasad et al.
2011/0164186	A1	7/2011	Sadek
2011/0182174	A1	7/2011	Pi et al.
2011/0206155	A1	8/2011	Fujimura et al.
2011/0235514	A1	9/2011	Huang et al.
2012/0058777	A1	3/2012	Nguyen et al.
2012/0063472	A1	3/2012	Le Pallec et al.
2012/0108284	A1	5/2012	Patel et al.
2012/0122477	A1	5/2012	Sadek et al.
2012/0135724	A1	5/2012	Lewis et al.
2012/0184222	A1	7/2012	Seok
2012/0213086	A1	8/2012	Matsuura et al.
2013/0089041	A1	4/2013	Negus et al.
2013/0089042	A1	4/2013	Negus et al.
2013/0089083	A1	4/2013	Negus et al.
2013/0100897	A1	4/2013	Negus
2013/0272217	A1	10/2013	Negus et al.
2013/0273974	A1	10/2013	Lea et al.
2013/0288593	A1	10/2013	Norin et al.
2013/0293419	A1	11/2013	Negus et al.
2014/0184455	A1	7/2014	Lea et al.
2014/0329562	A1	11/2014	Proctor, Jr. et al.
2015/0016362	A1	1/2015	Negus et al.
2015/0016561	A1	1/2015	Negus et al.
2015/0049720	A1	2/2015	Negus et al.
2015/0084826	A1	3/2015	Lea et al.
2015/0156777	A1	6/2015	Negus et al.

FOREIGN PATENT DOCUMENTS

EP	2839693	2/2015
EP	2843863	A1 3/2015
WO	2008007375	A2 1/2008
WO	2011/030960	A1 3/2011
WO	2011/080299	A2 7/2011
WO	WO 2013025413	A1 2/2013
WO	WO 2013055947	A1 4/2013
WO	WO 2013158237	A2 10/2013
WO	WO 2014040083	A2 3/2014

OTHER PUBLICATIONS

“Advances in Backhaul Synchronization—Maximizing ROI,” Application Brief, Symmetricom Inc., 2008, 6 pages.

“Carrier Ethernet Services Overview,” Presentation, MEF, Aug. 2008, 32 pages.

“Clock Distribution and Synchronization over Ethernet: IEEE1588v2 and SyncE,” Presentation, VITESSE, Sep. 2009, 9 pages.

“Clock Synchronization in Carrier Ethernet Networks—Synchronous Ethernet and 1588v2 on Juniper Networks MX Series 3D Universal Edge Routers,” White Paper, Juniper Networks, Inc., 2010, 11 pages.

“DAN2400-PTP—Open SoC Platform for Sub-6GHz Wireless Point-to-Point Systems,” DesignArt Networks, Jul. 2008, 2 pages.

“Doubling Capacity in Wireless Channels,” Provigent Inc., 2008, 3 pages.

“Evolving Microwave Mobile Backhaul for Next-Generation Networks,” White Paper, NEC Corp., Feb. 2008, 4 pages.

“GainSpan GS1011M Low Power Wi-Fi® Module Family,” Product Brief—Preliminary, GainSpan Corp., 2009, 2 pages.

“HMC536MS8G / 536MS8GE—GaAs MMIC Positive Control T/R Switch, DC—6 GHz,” Data Sheet, Switches—SMT, Hittite Microwave Corp., 2008, pp. 10.295-10.299.

“IEEE 1588v2 (Precision Time Protocol) Testing,” Application Note, IXIA, Dec. 2009.

“Information Technology—Open Systems Interconnection—Basic Reference Model: The Basic Model,” International Standard, ISO/IEC 7498-1, Second Edition Nov. 15, 1994, Corrected and Reprinted Jun. 15, 1996, 68 pages.

“MGA-21108—Broadband Fully Integrated Matched Low-Noise Amplifier MMIC,” Data Sheet, Avago Technologies, Aug. 2009, 21 pages.

“MGA-22003—2.3-2.7 GHz 3x3mm WiMAX/WiBro and WiFi Linear Amplifier Module,” Data Sheet, Avago Technologies, Mar. 2010, 15 pages.

“MGA-23003—3.3-3.8 GHz 3x3mm WiMAX Power Amplifier,” Data Sheet, Avago Technologies, Mar. 2010, 16 pages.

“MGA-25203—5.1-5.9GHz 3x3mm WiFi and WiMAX Power Amplifier,” Data Sheet, Avago Technologies, Mar. 2010, 14 pages.

“MGA-43328—High Linearity Wireless Data Power Amplifier for 2.5 to 2.7 GHz Applications,” Application Note, Avago Technologies, Apr. 2010, 10 pages.

“MGA-43328—(2.5-2.7) GHz 29dBm High Linearity Wireless Data Power Amplifier,” Data Sheet, Avago Technologies, Mar. 2010, 19 pages.

“MGA-645T6—Low Noise Amplifier with Bypass/Shutdown Mode in Low Profile Package,” Data Sheet, Avago Technologies, Nov. 2009, 16 pages.

“MGA-655T6—Low Noise Amplifier with Bypass Mode in Low Profile Package,” Data Sheet, Avago Technologies, Nov. 2009, 14 pages.

“MGA-675T6—Low Noise Amplifier with Shutdown Mode in Low Profile Package for 4.9-6 GHz Application,” Data Sheet, Avago Technologies, Nov. 2009, 16 pages.

“Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications—Amendment 5: Enhancements for Higher Throughput,” IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements, IEEE Computer Society, IEEE Std 802.11, Sep. 2009, 536 pages.

“Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific requirements, IEEE Computer Society, IEEE Std 802.11, Jun. 2007, 1233 pages.

“Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems—Amendment 2: Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands and Corrigendum 1,” IEEE Standard for Local and Metropolitan Area Networks, IEEE Computer Society and the IEEE Microwave Theory and Techniques Society, IEEE Std 802.16e-2005 and IEEE Std 802.16-2004/Cor1-2005, Feb. 2006, 864 pages.

“Part 16: Air Interface for Fixed Broadband Wireless Access Systems,” IEEE Standard for Local and Metropolitan Area Networks, IEEE Std 802.16, IEEE Computer Society and the IEEE Microwave Theory and Techniques Society, Oct. 2004, 895 pages.

“PC203-PC203-10—Basestation PHY Processor,” Wireless Communications Processors, Product Brief, picoChip Flexible Wireless, 2008, 4 pages.

“SC-FDMA—Single Carrier FDMA in LTE,” Data Sheet, IXIA Corp., Nov. 2009, 16 pages.

“Spectrum Sharing: The Promise and the Reality,” RYSAVY Research, LLC, White Paper, Jul. 2012, pp. 1-12, available at <http://www.rysavvy.com>.

“Technical Report: Electromagnetic Compatibility and Radio Spectrum Matters (ERM); System Reference Document (SRdoc); Mobile Broadband Services in the 2 300 MHz-2 400 MHz frequency band under Licensed Shared Access Regime”, ETSI TR 103 113 V1.1.1 (Jul. 2013), European Telecommunications Standards Institute, France, 2013, pp. 1-37.

“Understanding MPLS-TP and Its Benefits,” White Paper, Cisco Systems Inc., 2009, 6 pages.

“WiMAX/WiBro RF MxFE Transceiver, AD9352,” Data Sheet, Analog Devices, 2007.

“WiMAX RF MxFE Transceiver, AD9352-5,” Data Sheet, Analog Devices, 2008.

“WiMAX RF MxFE Transceiver, AD9353,” Data Sheet, Analog Devices, 2007.

“WiMAX/WiBro RF MxFE MISO Transceiver, AD9354,” Data Sheet, Analog Devices, 2008-2009.

(56)

References Cited**OTHER PUBLICATIONS**

- "WiMAX/WiBro RF MxFE MISO Transceiver, AD9355," Data Sheet, Analog Devices, 2008-2009.
- "WiMAX/BWA/WiBro/LTE RF MxFE 2x2 MIMO Transceiver, AD9356," Data Sheet, Analog Devices, 2010.
- "WiMAX/BWA/LTE RF MxFE 2x2 MIMO Transceiver, AD9357," Data Sheet, Analog Devices, 2010.
- Baum, D.S. et al., "An Interim Channel Model for Beyond-3G Systems—Extending the 3GPP Spatial Channel Model (SCM)," IEEE Vehicular Technology Conference, vol. 5, 2005, pp. 3132-3136.
- Beller, D. et al., "MPLS-TP—The New Technology for Packet Transport Networks," Alcatel-Lucent Deutschland AG, Stuttgart, 2009, 11 pages.
- Chundury, R., "Mobile Broadband Backhaul: Addressing the Challenge," Planning Backhaul Networks, Ericsson Review, 2008, pp. 4-9.
- Conrat, J.-M. et al., "Typical MIMO Propagation Channels in Urban Macrocells at 2 GHz," EURASIP Journal on Wireless Communications and Networking, vol. 2007, Iss. 2, Jan. 2007, 9 pages.
- Coon, J.P. et al., "Adaptive Frequency-Domain Equalization for Single-Carrier Multiple-Input Multiple-Output Wireless Transmissions," IEEE Transactions on Signal Processing, vol. 53, No. 8, Aug. 2005, pp. 3247-3256.
- Coon, J.P. et al., "An Investigation of MIMO Single-Carrier Frequency-Domain MMSE Equalization," Centre for Communications Research, University of Bristol, Sep. 2002, 4 pages.
- Durgin, G.D., "Wideband Measurements of Angle and Delay Dispersion for Outdoor and Indoor Peer-to-Peer Radio Channels at 1920 MHz," IEEE Transactions on Antennas and Propagation, vol. 51, No. 5, May 2003, pp. 936-944.
- Falconer, D.D. et al., "Broadband Wireless Using Single Carrier and Frequency Domain Equalization," Invited Overview Paper for WPMC '02, Honolulu, Oct. 2002, 10 pages.
- Falconer, D.D. et al., "Frequency Domain Equalization for Single-Carrier Broadband Wireless Systems," Wideband Wireless Access Technologies to Broadband Internet, IEEE Communications Magazine, Apr. 2002, pp. 58-66.
- Frost, D. et al., "MPLS Transport Profile Data Plane Architecture," Internet Engineering Task Force, RFC 5960, Aug. 2010, 16 pages.
- Gao, S.C. et al., "Dual-polarized slot coupled planar antenna with wide bandwidth," IEEE Trans. Antennas and Propagation, vol. 51, No. 3, pp. 441-448, 2003.
- Garner, G.M., "IEEE 1588 Version 2," ISPCS Ann Arbor '08, Sep. 2008, 89 pages.
- Hentschel, T. et al., "The Digital Front-End—Bridge Between RF and Baseband-Processing," Software Defined Radio: Enabling Technologies by Walter Tuttlebee, Chapter 6, Jun. 2002, 58 pages.
- Lashkarian, N., "Reconfigurable Digital Front-end Architecture for Wireless Base-Station Transmitters: Analysis, Design and FPGA Implementation," Seminar Presentation, University of California, Berkeley, Apr. 2006, 86 pages.
- Nathanzon, G. et al., "Cost-Effective Backhaul Alternatives for WiMAX Networks: Challenges & Solutions," Provigent Inc., Jun. 2008, 16 pages.
- Padhi, S.K. et al., "A Dual Polarized Aperture Coupled Circular Patch Antenna Using a C-Shaped Coupling Slot," IEEE Transactions on Antennas and Propagation, vol. 51, No. 12, Dec. 2003, pp. 3295-3298.
- Pancaldi, F. et al., "Single-Carrier Frequency Domain Equalization—A Focus on Wireless Applications," IEEE Signal Processing Magazine, Sep. 2008, 22 pages.
- Pozar, D.M. et al., "Improved coupling for aperture-coupled microstrip antennas," Electron. Lett., vol. 27, No. 13, pp. 1129-1131, 1991.
- Pozar, D.M., "A microstrip antenna aperture-coupled to a microstripline," Electron. Lett., vol. 21, No. 2, pp. 49-50, 1985.
- Sharony, J., "Introduction to Wireless MIMO—Theory and Applications," CEWIT—Center of Excellence in Wireless and Informational Technology, Stony Brook University, IEEE LI, Nov. 15, 2006, 63 pages.
- Soffer, R., "Microwave Wireless Backhaul for LTE Networks—Meeting the Rapidly-Increasing Demand for Mobile Data," Provigent Inc., Jan. 2010, 18 pages.
- Soffer, R., "ProviBand—Next Generation Wireless Backhaul Technology," Technical White Paper, Provigent Inc., Jan. 2006, 13 pages.
- Stuber, G.L. et al., "Broadband MIMO-OFDM Wireless Communications," Invited Paper, Proceedings of the IEEE, vol. 92, No. 2, Feb. 2004, pp. 271-294.
- Tubbaix, J., "OFDM Versus Single Carrier with Cyclic Prefix: A System-based Comparison for Binary Modulation," IMEC, Belgium, Jan. 2011, 4 pages.
- "TV Whitespaces" reuse: A detailed description of the program is provided in FCC order FCC-10-174A1, and the rules for unlicensed devices that operate in the TV bands are set forth in 47 C.F.R. §§ 15.701-717. See Title 47—Telecommunication; Chapter I—Federal Communications Commission; Subchapter A—General, Part 15—Radio Frequency Devices, Subpart H—Television Band Devices (Available online at <http://www.ecfr.gov/cgi-bin/text-idx?c=ecfr&SID=30f46f0753577b10de41d650c7adf941&rgn=div6&view=text&node=47:1.0.1.1.16.8&idno=47:1.0.1.1.16.8&idno=47>, retrieved Feb. 25, 2015).
- Notice of Allowance in U.S. Appl. No. 13/212,036, dated Apr. 3, 2012.
- Office Action in U.S. Appl. No. 13/271,051, dated Jun. 14, 2013.
- Amendment and Response to Office Action in U.S. Appl. No. 13/271,051, filed Dec. 16, 2013.
- Notice of Allowance in U.S. Appl. No. 13/271,051, dated Feb. 3, 2014.
- Office Action in U.S. Appl. No. 13/371,346, dated Sep. 26, 2012.
- Amendment and Response to Office Action in U.S. Appl. No. 13/371,346, filed Dec. 21, 2012.
- Applicant Initiated Interview Summary in U.S. Appl. No. 13/371,346, filed Jan. 10, 2013.
- Supplemental Amendment and Response to Office Action in U.S. Appl. No. 13/371,346, filed Jan. 25, 2013.
- Notice of Allowance in U.S. Appl. No. 13/371,346, dated Apr. 2, 2013.
- Office Action in U.S. Appl. No. 13/371,366, dated Apr. 19, 2012.
- Amendment and Response to Office Action in U.S. Appl. No. 13/371,366, dated Jul. 19, 2012.
- Notice of Allowance in U.S. Appl. No. 13/371,366, dated Sep. 21, 2012.
- Office Action in U.S. Appl. No. 13/415,778, mailed Apr. 26, 2012.
- Applicant Initiated Interview Summary in U.S. Appl. No. 13/415,778, filed Jul. 5, 2012.
- Amendment and Response to Office Action in U.S. Appl. No. 13/415,778, dated Jul. 26, 2012.
- Notice of Allowance in U.S. Appl. No. 13/415,778, mailed Sep. 17, 2012.
- Office Action in U.S. Appl. No. 13/448,294, dated on Jun. 21, 2012.
- Amendment and Response to Office Action in U.S. Appl. No. 13/448,294, dated Sep. 21, 2012.
- Notice of Allowance in U.S. Appl. No. 13/448,294, dated Oct. 23, 2012.
- Notice of Allowance in U.S. Appl. No. 13/536,927, dated Feb. 19, 2013.
- Notice of Allowance in U.S. Appl. No. 29/429,634 dated Dec. 5, 2013.
- Office Action in U.S. Appl. No. 13/632,961, dated May 6, 2014.
- Office Action in U.S. Appl. No. 13/632,993, dated Jan. 30, 2015.
- Notice of Allowance in U.S. Appl. No. 13/633,028, dated May 1, 2014.
- Office Action in U.S. Appl. No. 13/645,472, dated May 24, 2013.
- Amendment and Response to Office Action in U.S. Appl. No. 13/645,472, dated Nov. 25, 2013.
- Supplemental Amendment and Response to Office Action in U.S. Appl. No. 13/645,472, dated Jan. 27, 2014.
- Notice of Allowance in U.S. Appl. No. 13/645,472, dated Apr. 4, 2014.
- Notice of Allowance in U.S. Appl. No. 13/748,544, dated Aug. 15, 2014.
- Notice of Allowance in U.S. Appl. No. 13/898,429, dated Apr. 8, 2014.

(56)

References Cited

OTHER PUBLICATIONS

Office Action in U.S. Appl. No. 13/934,175, dated Oct. 15, 2014.
 Notice of Allowance in U.S. Appl. No. 14/197,158, dated Sep. 3, 2014.
 Examiner Initiated Interview Summary in U.S. Appl. No. 14/197,158, filed Sep. 3, 2014.
 Examiner Initiated Interview Summary in U.S. Appl. No. 14/199,734, dated Jun. 13, 2014.
 Notice of Allowance in U.S. Appl. No. 14/199,734, dated Jun. 13, 2014.
 Notice of Allowance in U.S. Appl. No. 14/336,958, dated Nov. 26, 2014.
 Preliminary Amendment and Response in U.S. Appl. No. 14/337,744, dated Aug. 20, 2014.
 Applicant Response and Amendment in U.S. Appl. No. 14/337,744, dated Oct. 6, 2014.
 Notice of Allowance in U.S. Appl. No. 14/337,744, dated Jan. 12, 2015.
 Office Action in U.S. Appl. No. 14/502,471, dated Jan. 5, 2015.
 Notice of Allowance in U.S. Appl. No. 14/498,959, dated Jan. 12, 2015.
 Office Action in U.S. Appl. No. 14/552,431, dated Mar. 12, 2015.
 International Search Report and Written Opinion for PCT Application No. PCT/US2012/049948 dated Oct. 22, 2012.

International Search Report and Written Opinion for PCT Application No. PCT/US2012/059797 dated Jan. 2, 2013.
 International Search Report and Written Opinion for PCT Application No. PCT/US2013/029731 dated May 13, 2013.
 International Search Report and Written Opinion in PCT Application No. PCT/US013/48782, dated Jul. 29, 2013.
 Extended European Search Report in European Patent Application No. 14177662.5, dated Feb. 3, 2015.
 U.S. Appl. No. 14/552,431, filed Nov. 24, 2014.
 U.S. Appl. No. 14/624,365, filed Feb. 17, 2015.
 Notice of Allowance in U.S. Appl. No. 13/934,175, dated May 28, 2015.
 U.S. Appl. No. 14/686,674, filed Apr. 14, 2015.
 U.S. Appl. No. 14/688,550, filed Apr. 16, 2015.
 Extended European Search Report in European Patent Application No. 12824072.8, dated Jul. 20, 2015.
 Office Action in U.S. Appl. No. 14/502,471, dated Jul. 22, 2015.
 Final Office Action in U.S. Appl. No. 14/552,431, dated Jul. 14, 2015.
 Office Action in U.S. Appl. No. 14/686,674, dated Jul. 23, 2015.
 Office Action in U.S. Appl. No. 14/688,550, dated Jul. 29, 2015.
 Amendment and Response to Office Action in U.S. Appl. No. 13/632,993, dated Jul. 27, 2015.
 Applicant Response and Amendment in U.S. Appl. No. 14/502,471, dated Jul. 2, 2015.
 Notice of Allowance in U.S. Appl. No. 14/552,431, dated Aug. 19, 2015.
 Partial European Search Report in European Patent Application No. 12839752.8, dated Jun. 16, 2015.




FIG. 1 (Prior Art)

200




FIG. 2 (Prior Art)

300

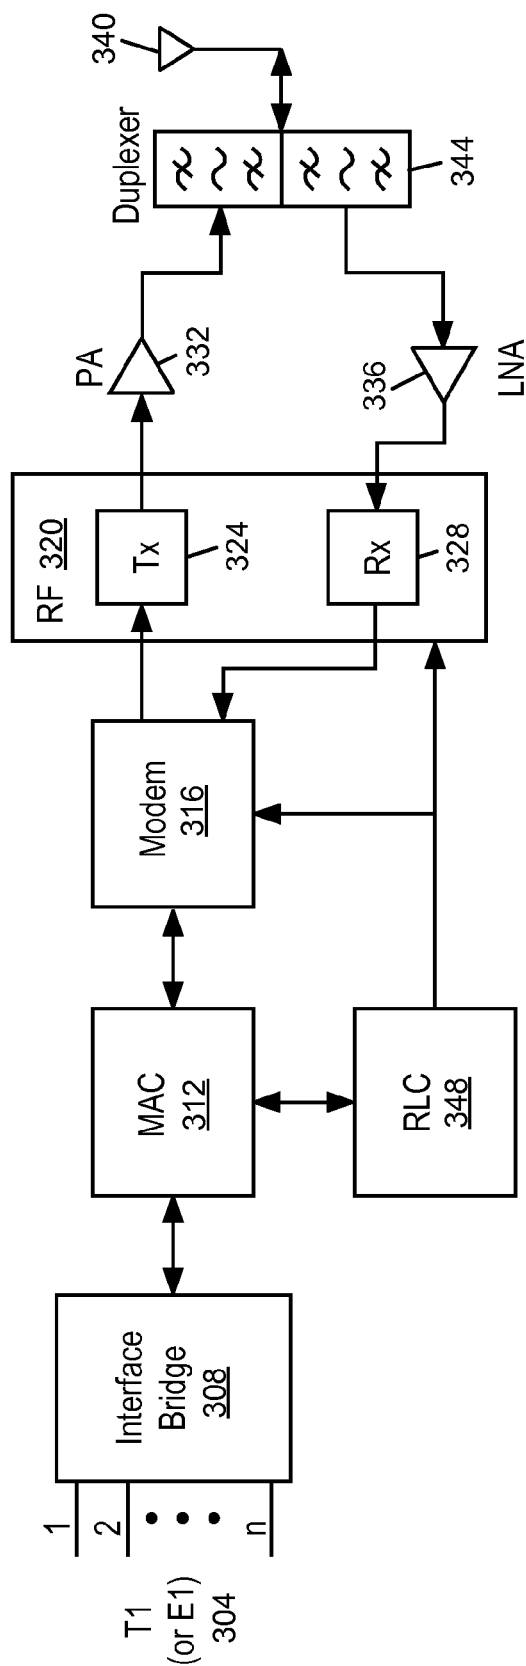


FIG. 3 (Prior Art)

400




FIG. 4




FIG. 5

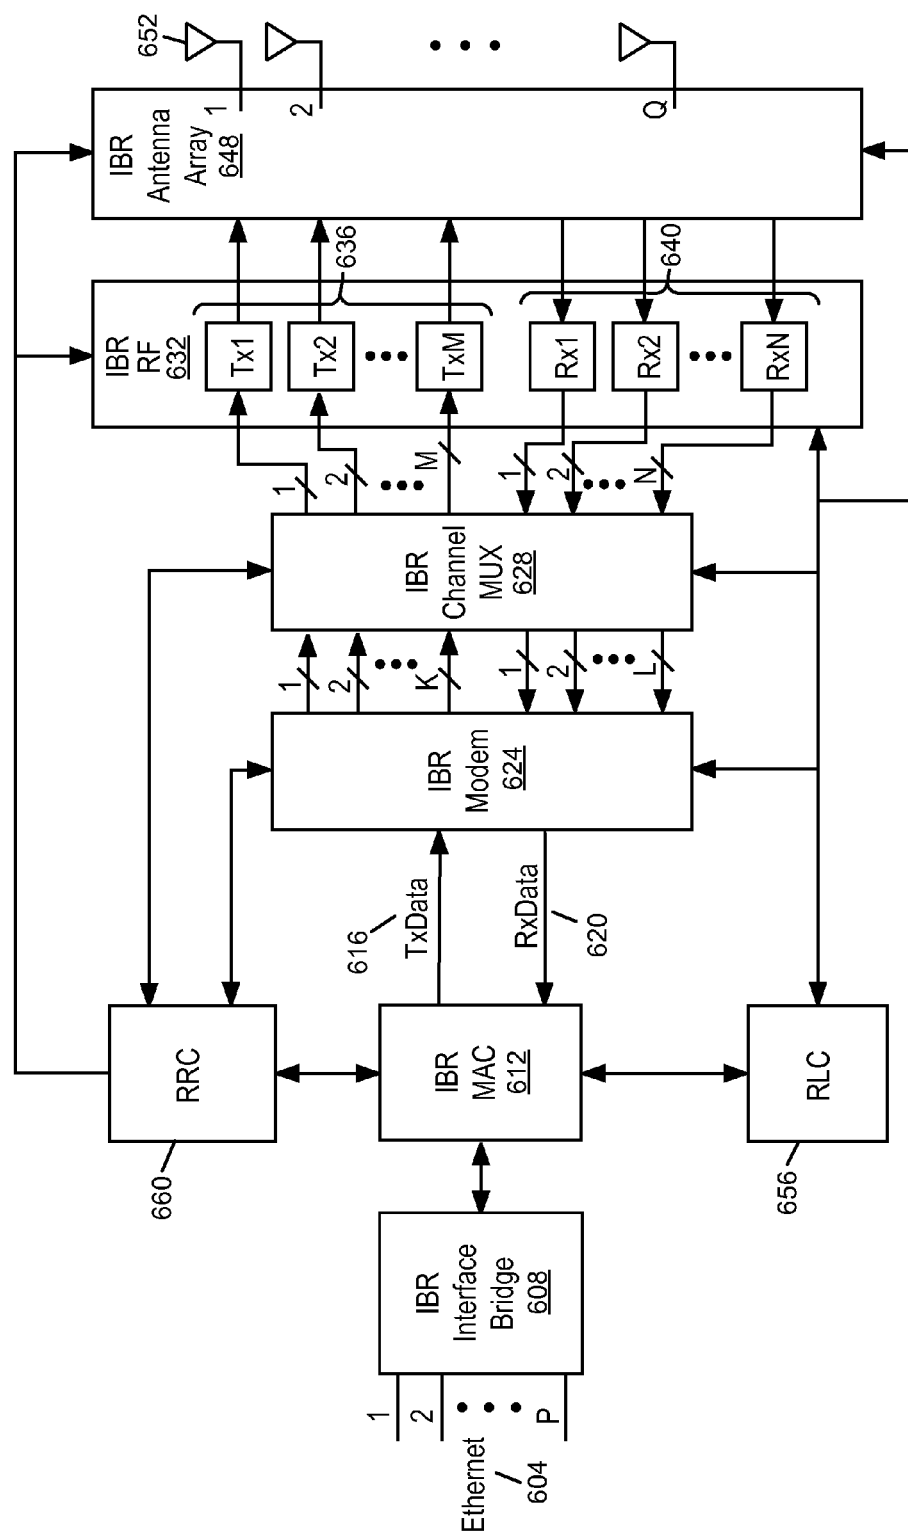


FIG. 6

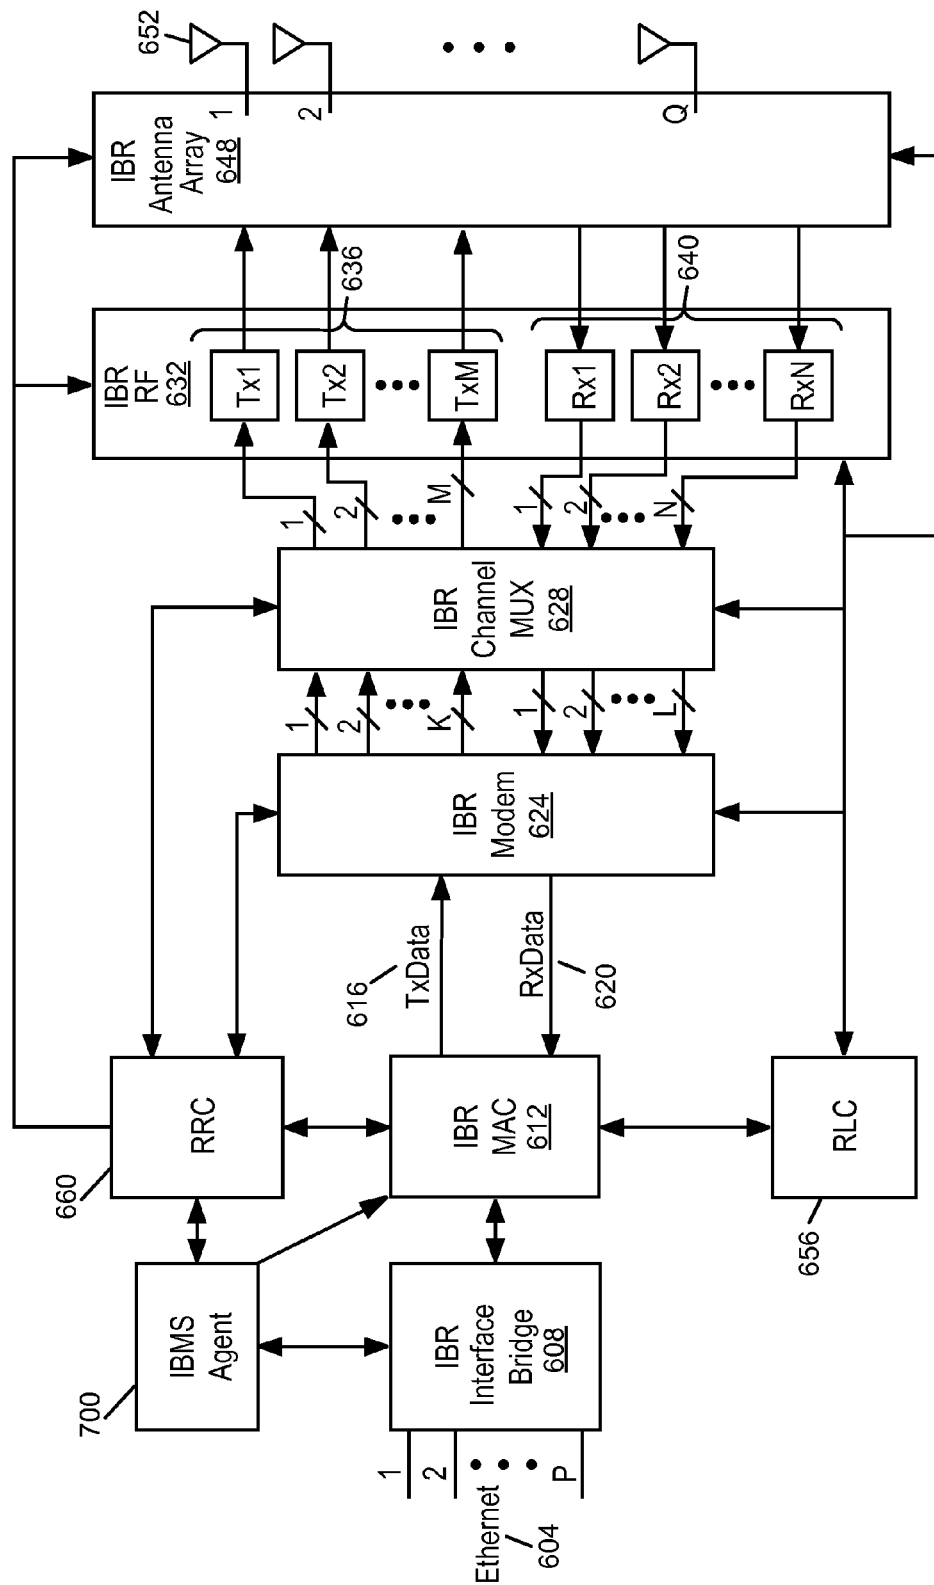


FIG. 7




FIG. 8

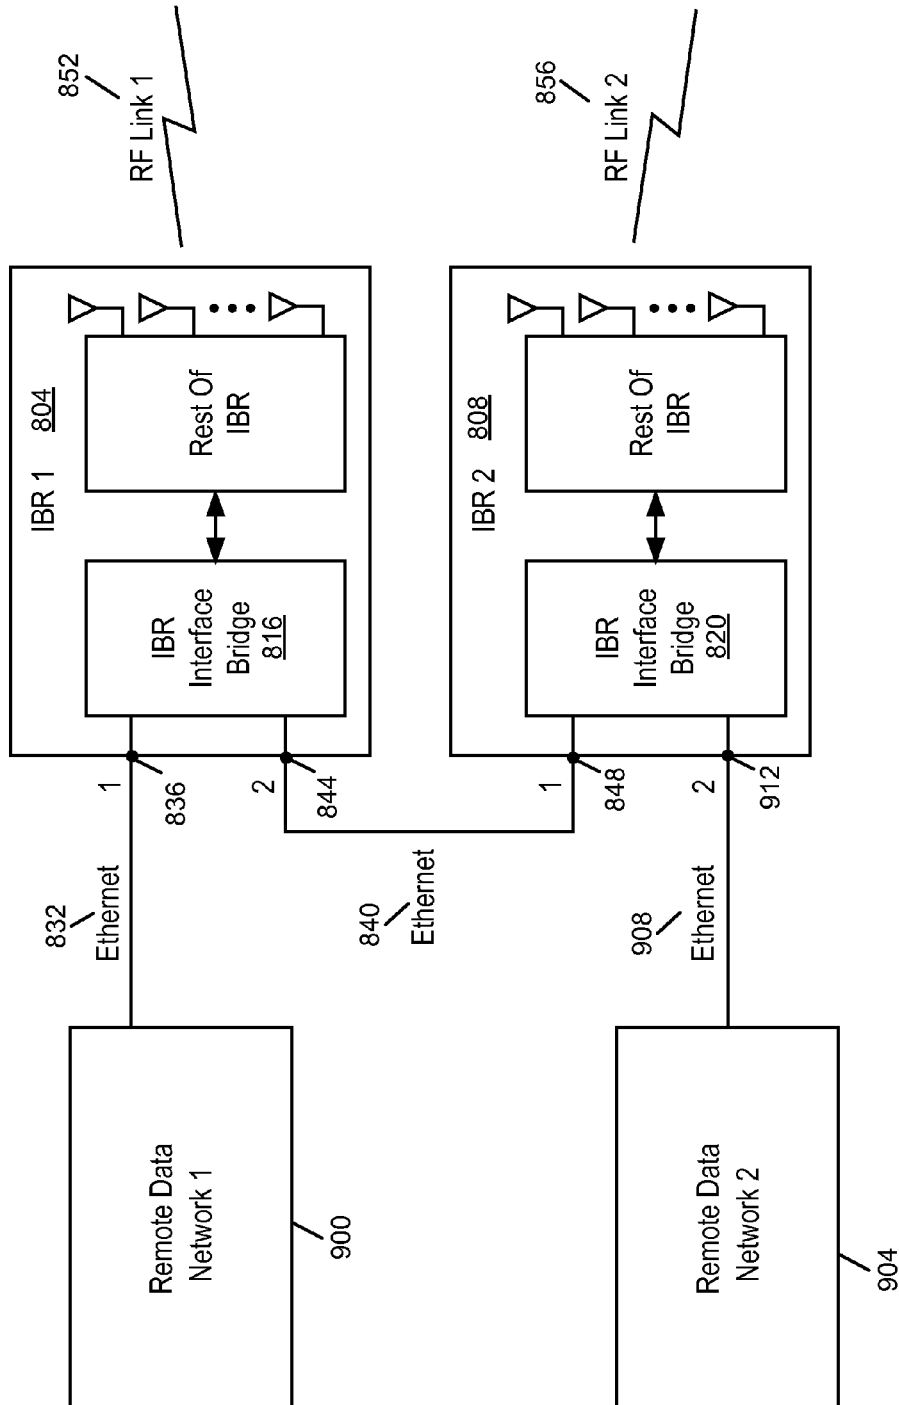


FIG. 9

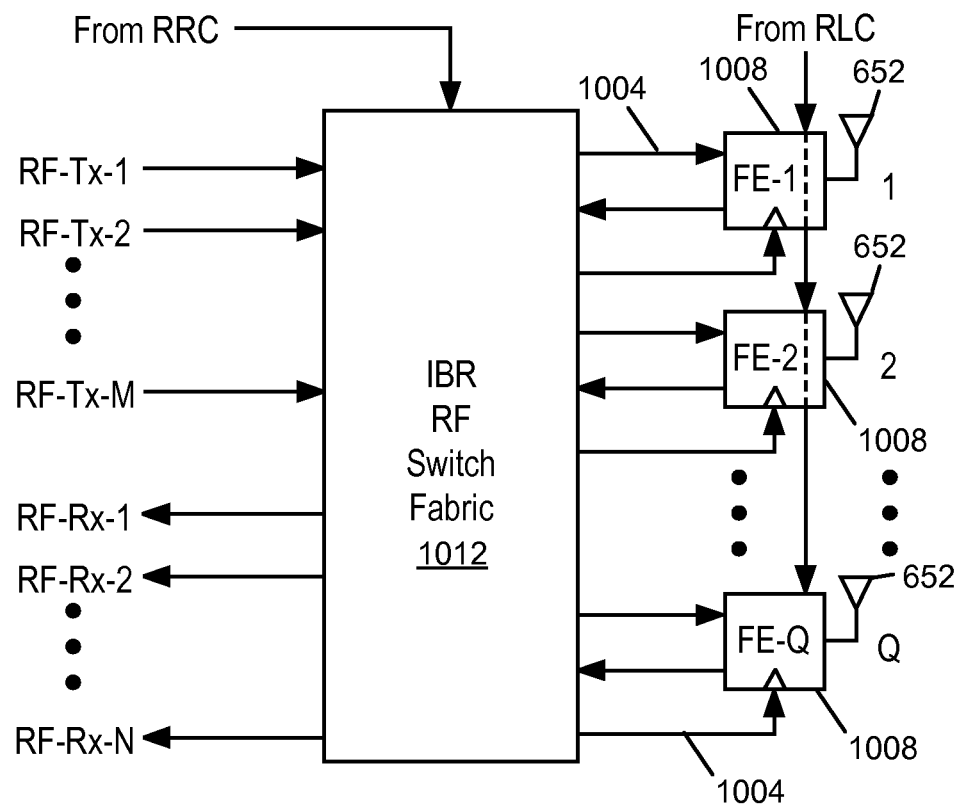
648

FIG. 10

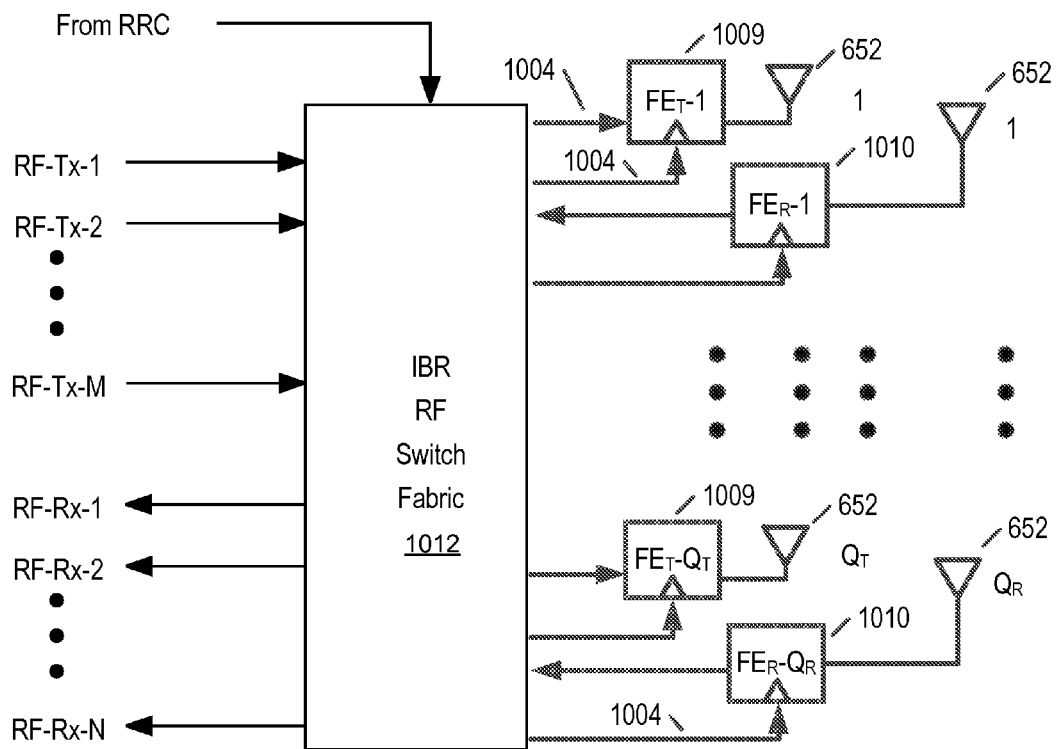


FIG. 10A




FIG. 11

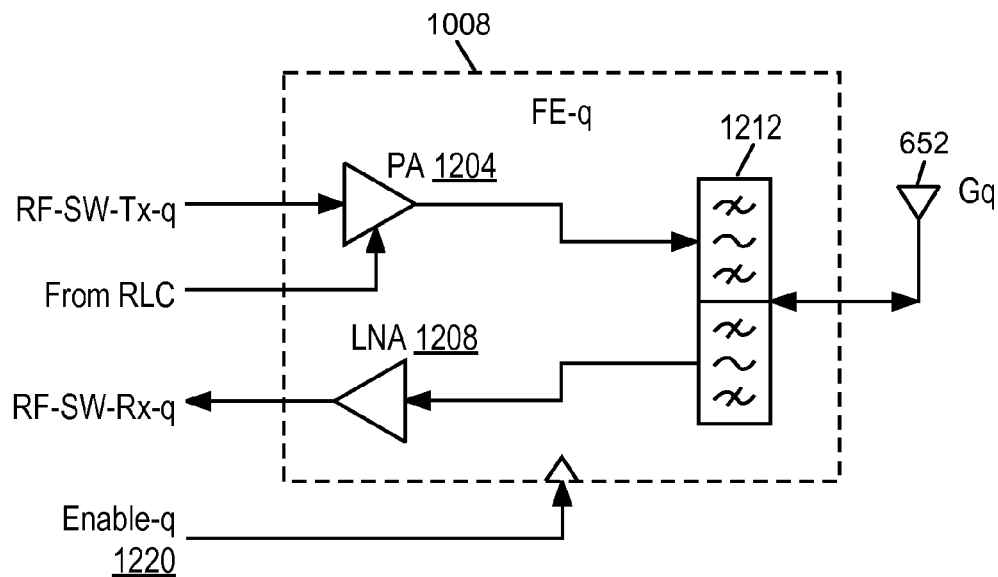


FIG. 12




Fig. 12A




Fig. 12B




FIG. 13




FIG. 14




FIG. 15




FIG. 16


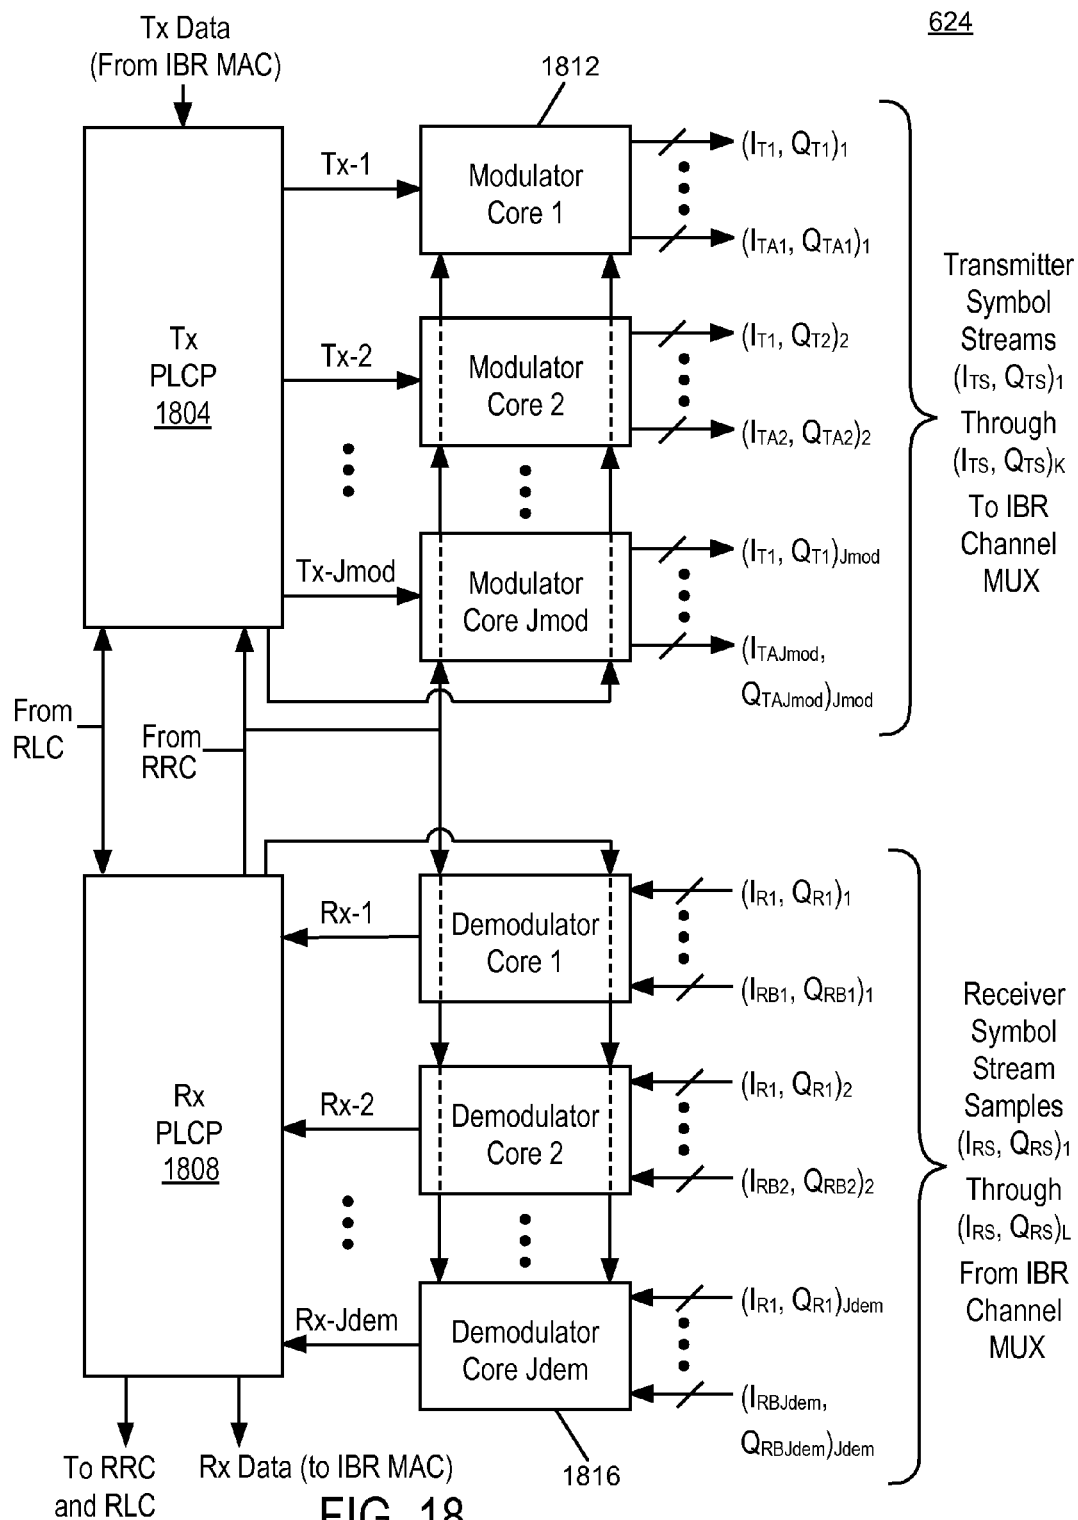


FIG. 17



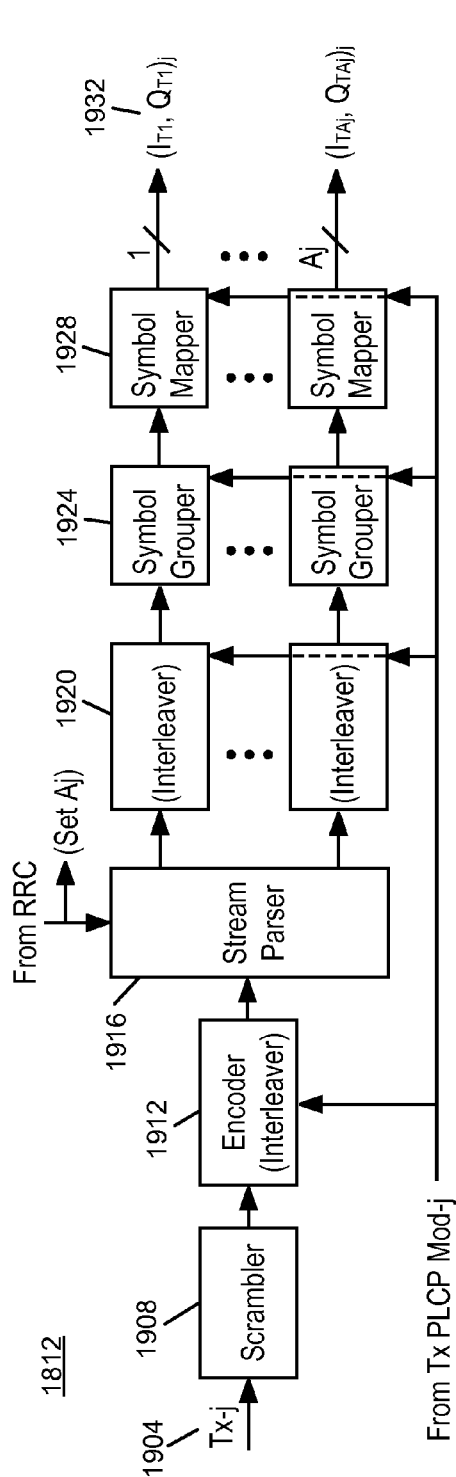


FIG. 19

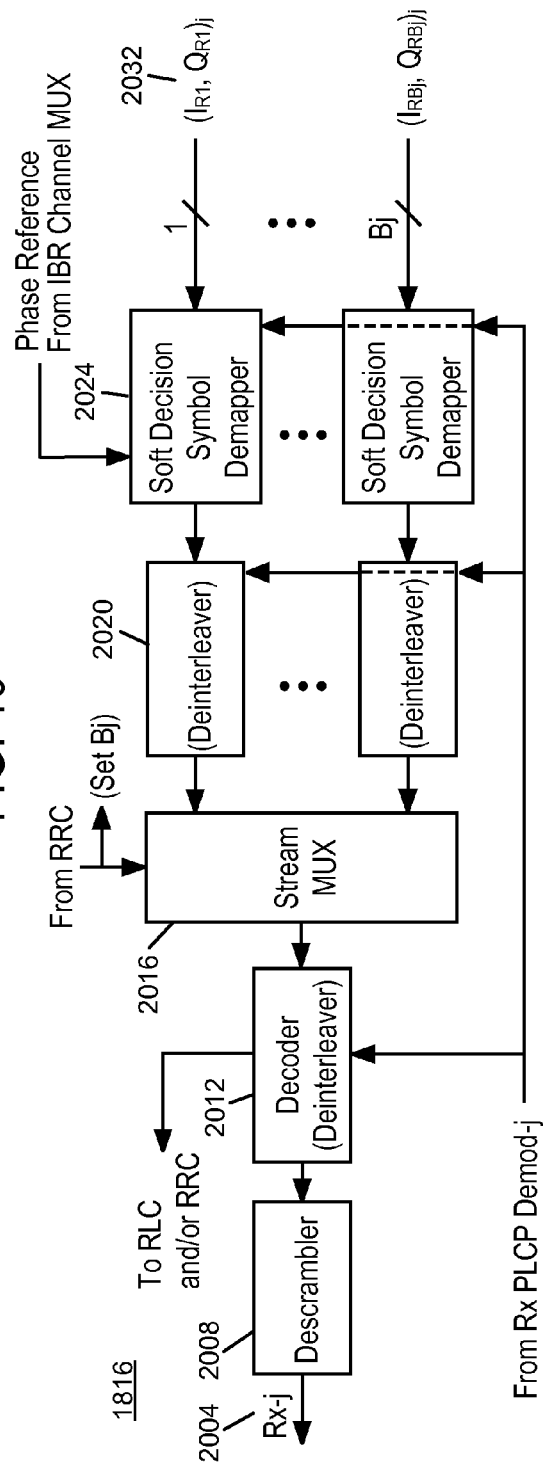
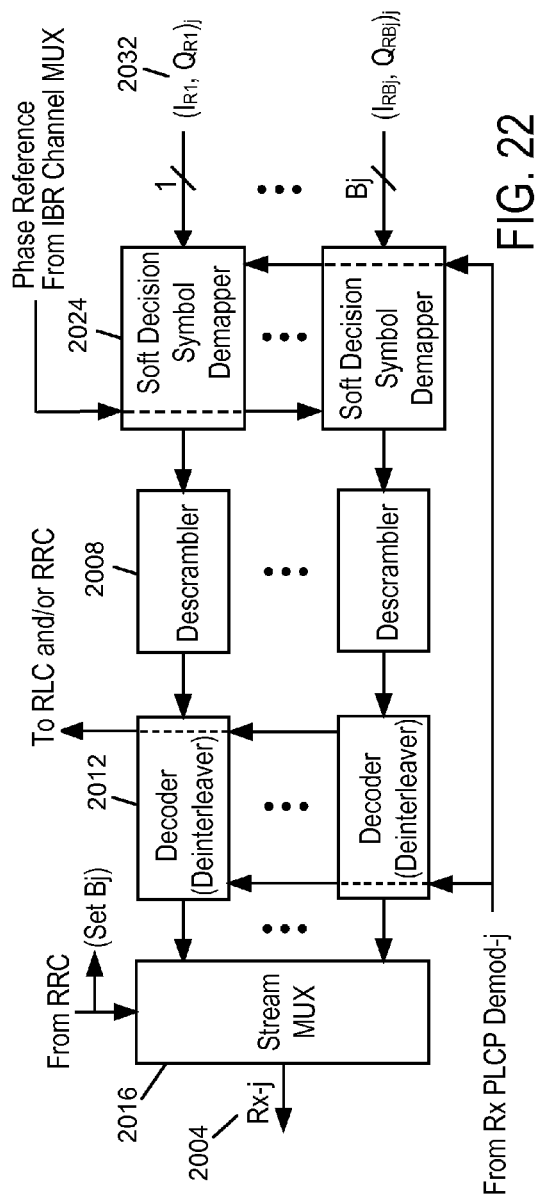



FIG. 20



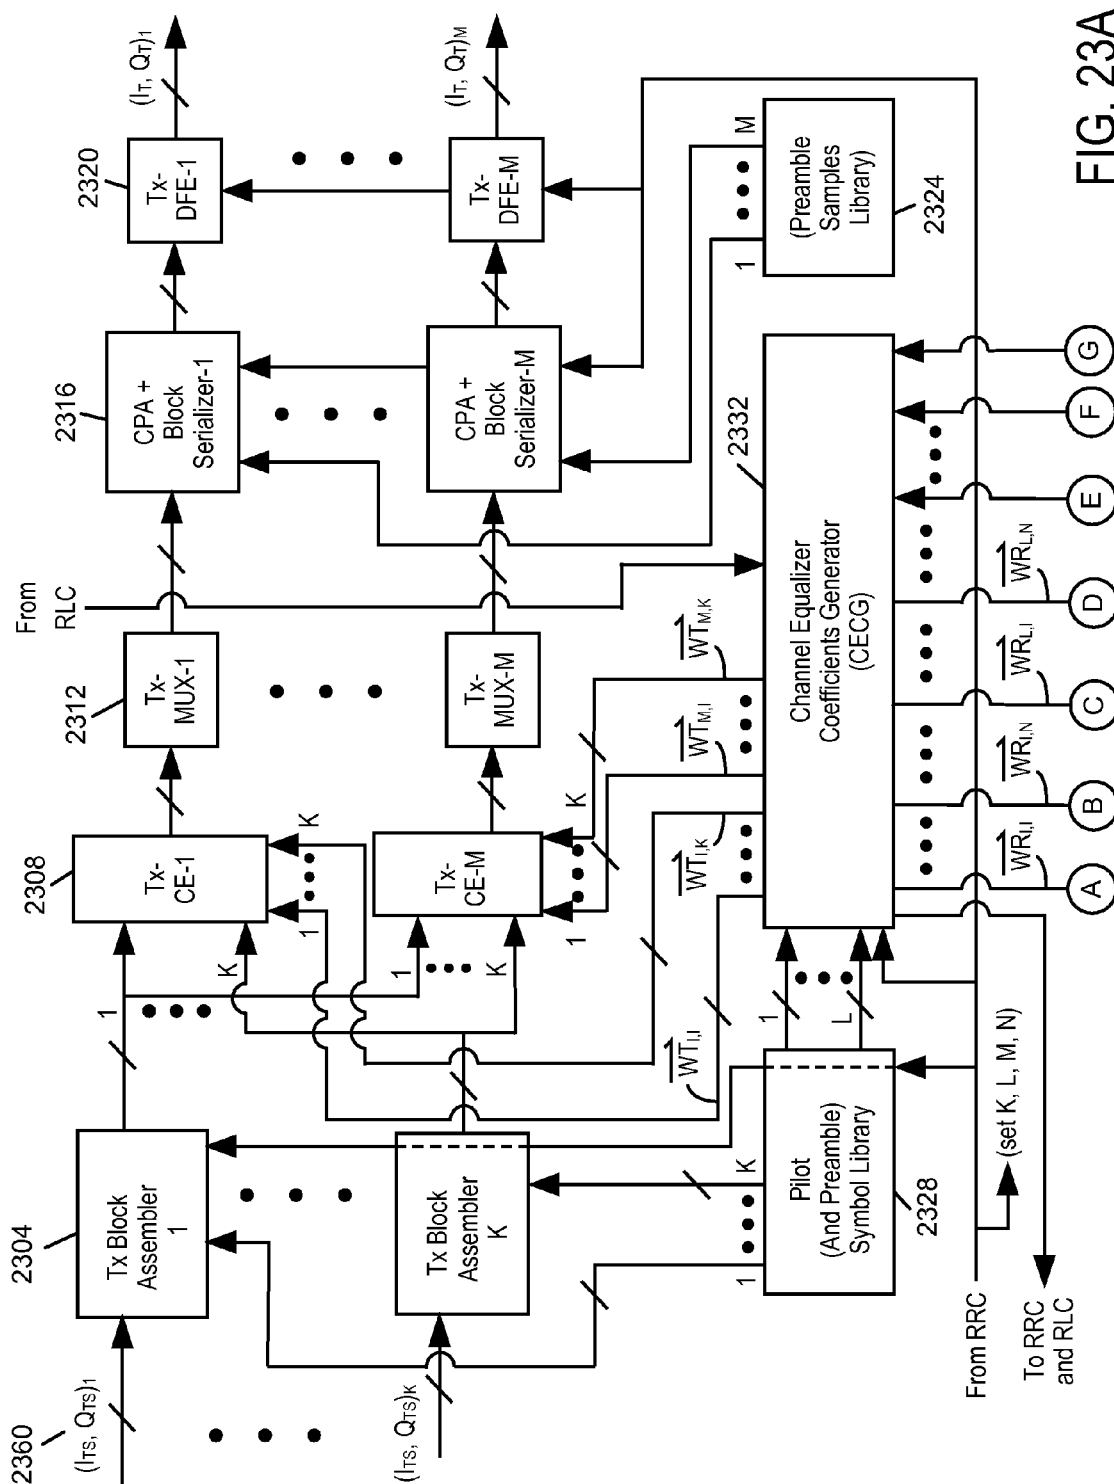


FIG. 23A

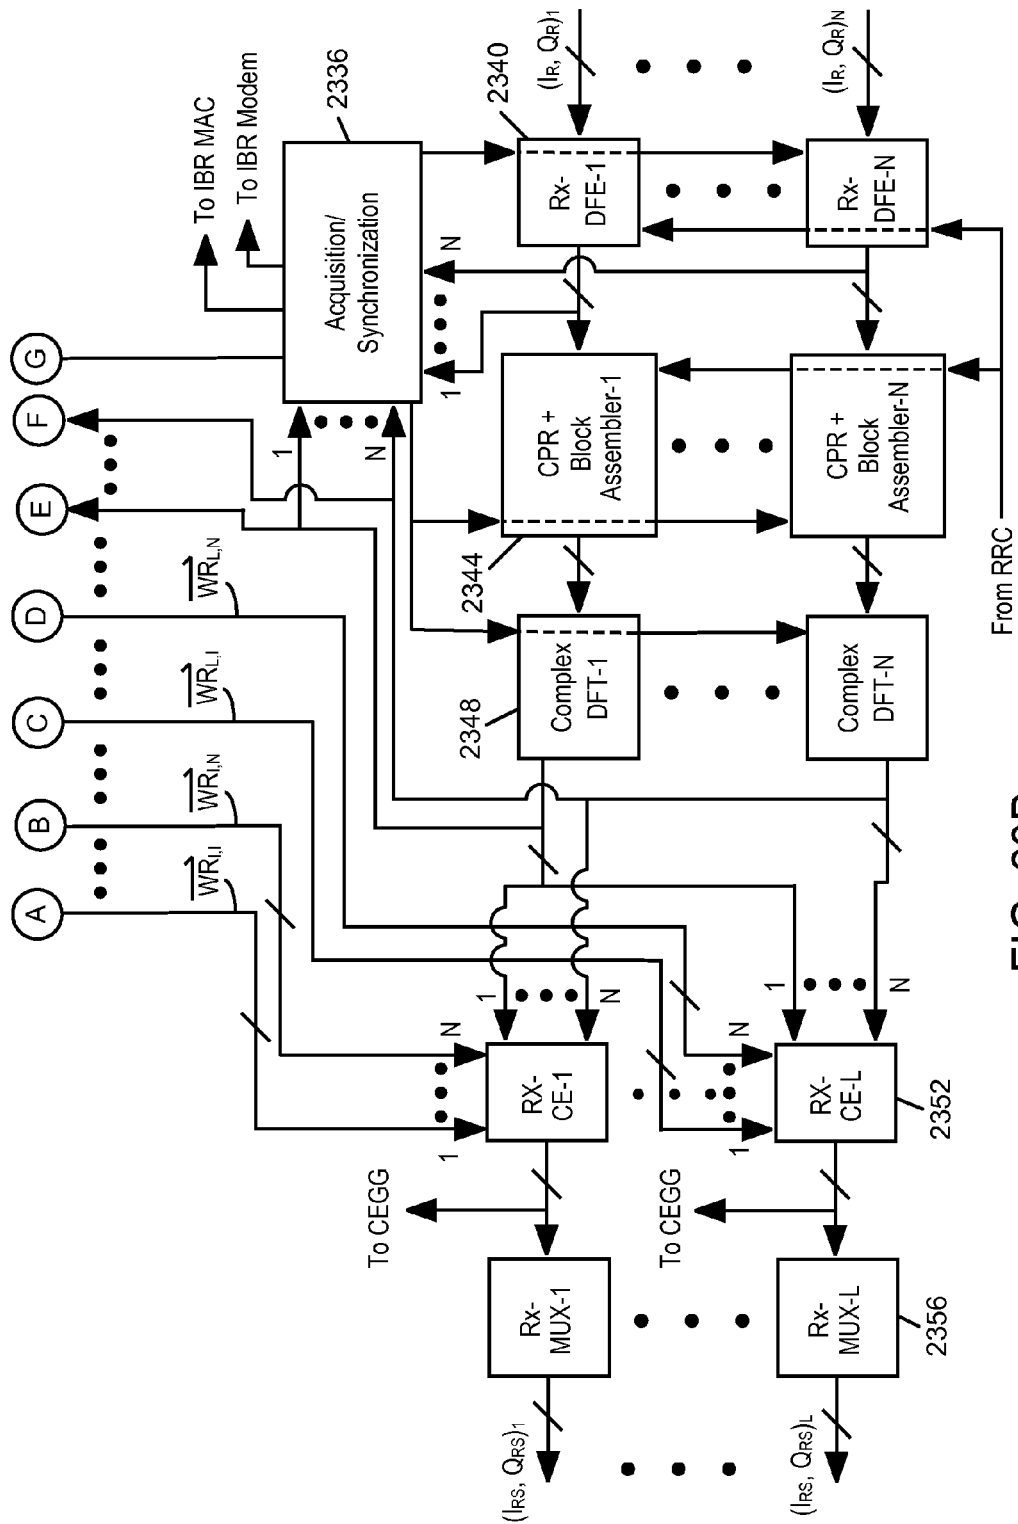


FIG. 23B

2308




FIG. 24

2352




FIG. 25

2308

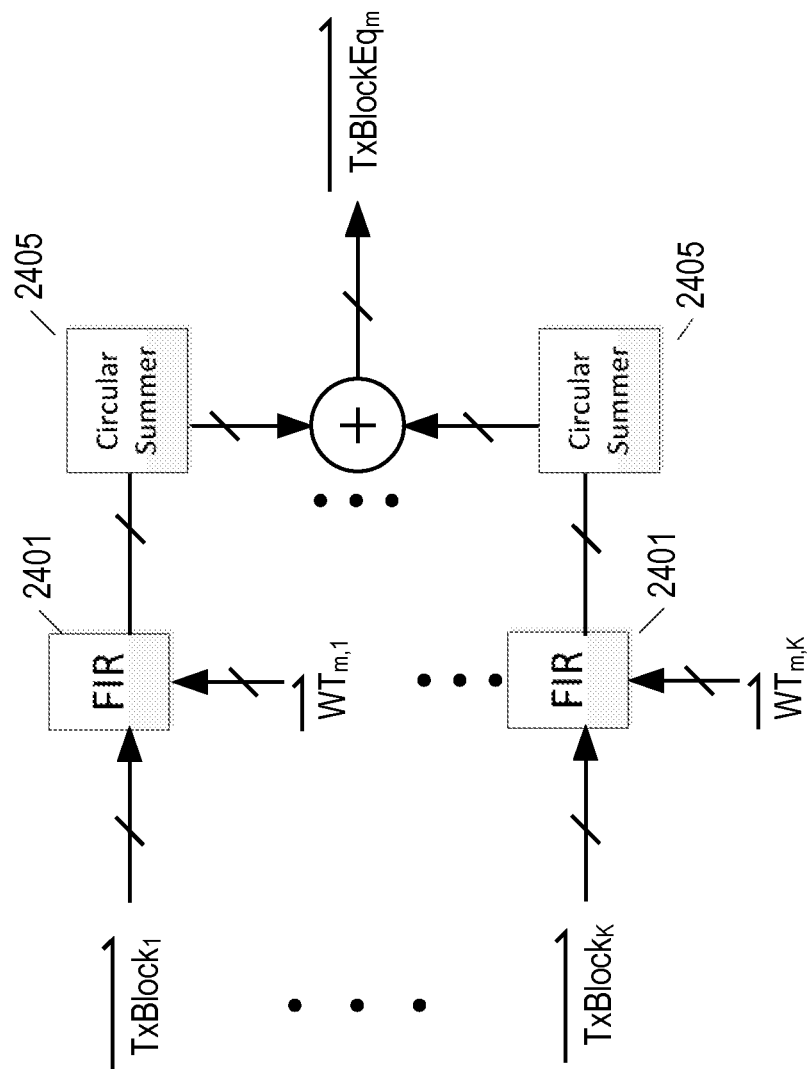


FIG. 24A




Fig. 24B




Fig. 24C




Fig. 24D




Fig. 24E




Fig. 24F




Fig. 24G




Fig. 24H




Fig. 24I




FIG. 26

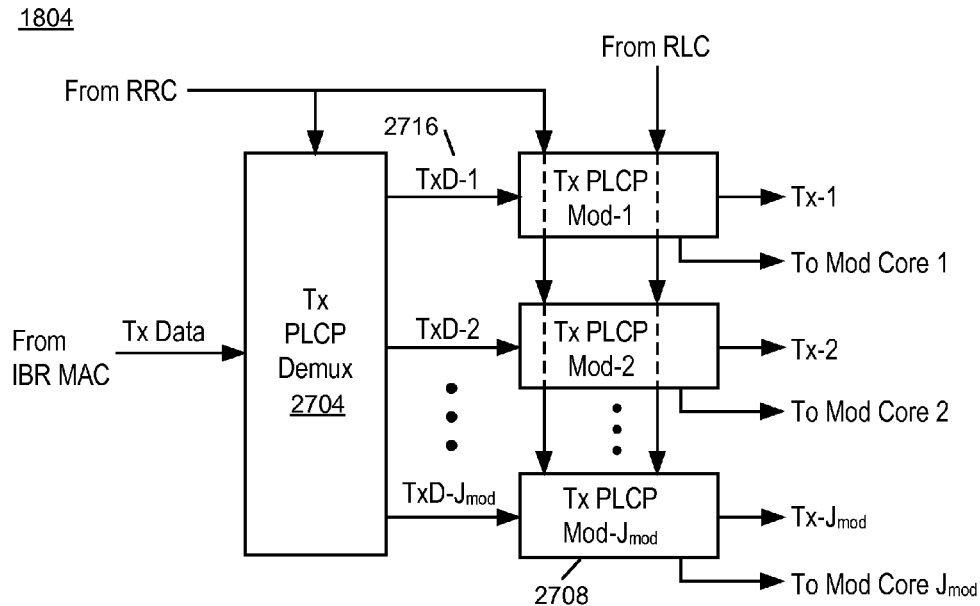


FIG. 27

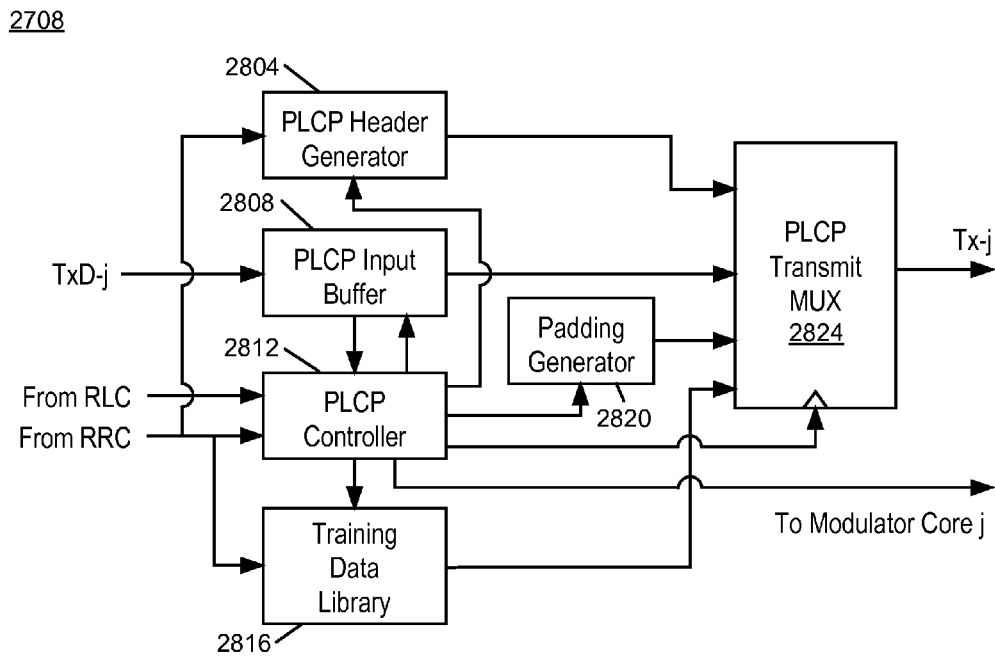


FIG. 28

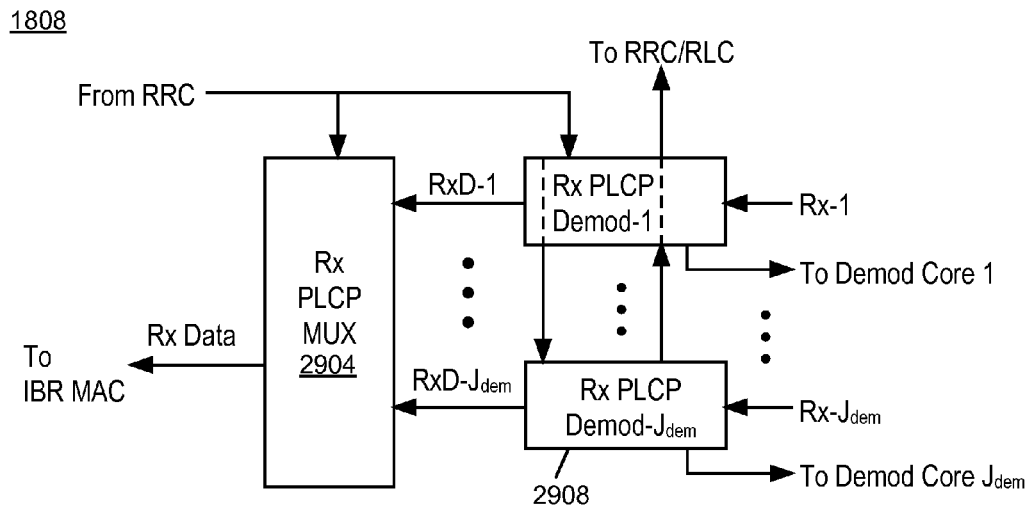


FIG. 29

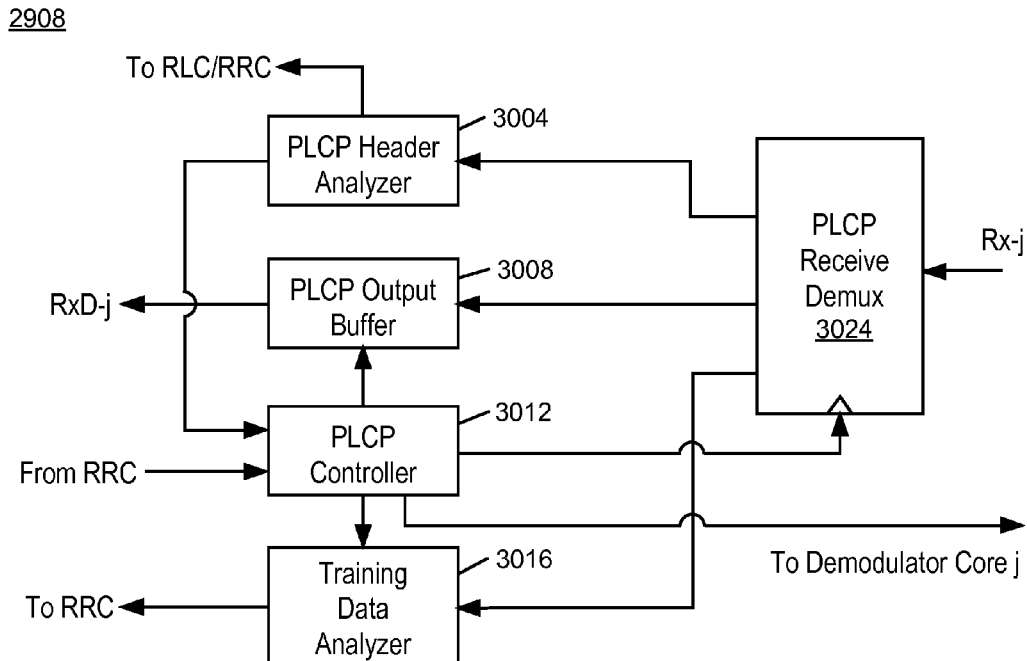


FIG. 30




FIG. 31




FIG. 32

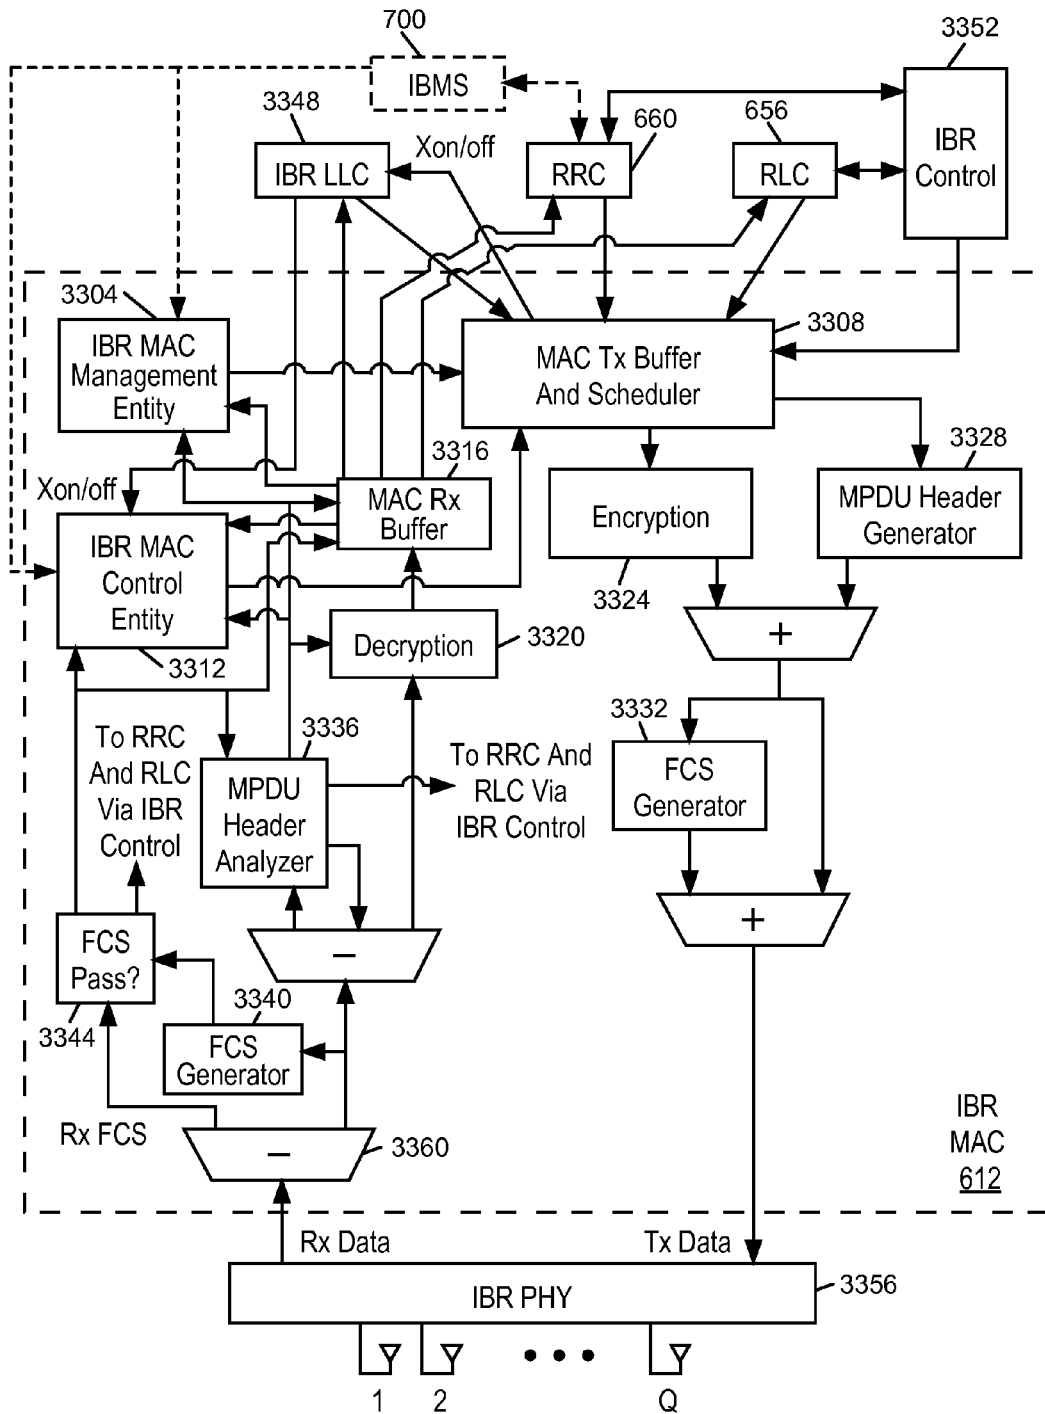


FIG. 33




FIG. 34




FIG. 35




FIG. 36




Fig. 37




Fig. 38




Fig. 39




Fig. 40




Fig. 41




Fig. 42




Fig. 43




Fig. 44A




Fig. 44B




Fig. 45A




Fig. 45B




Fig. 46A

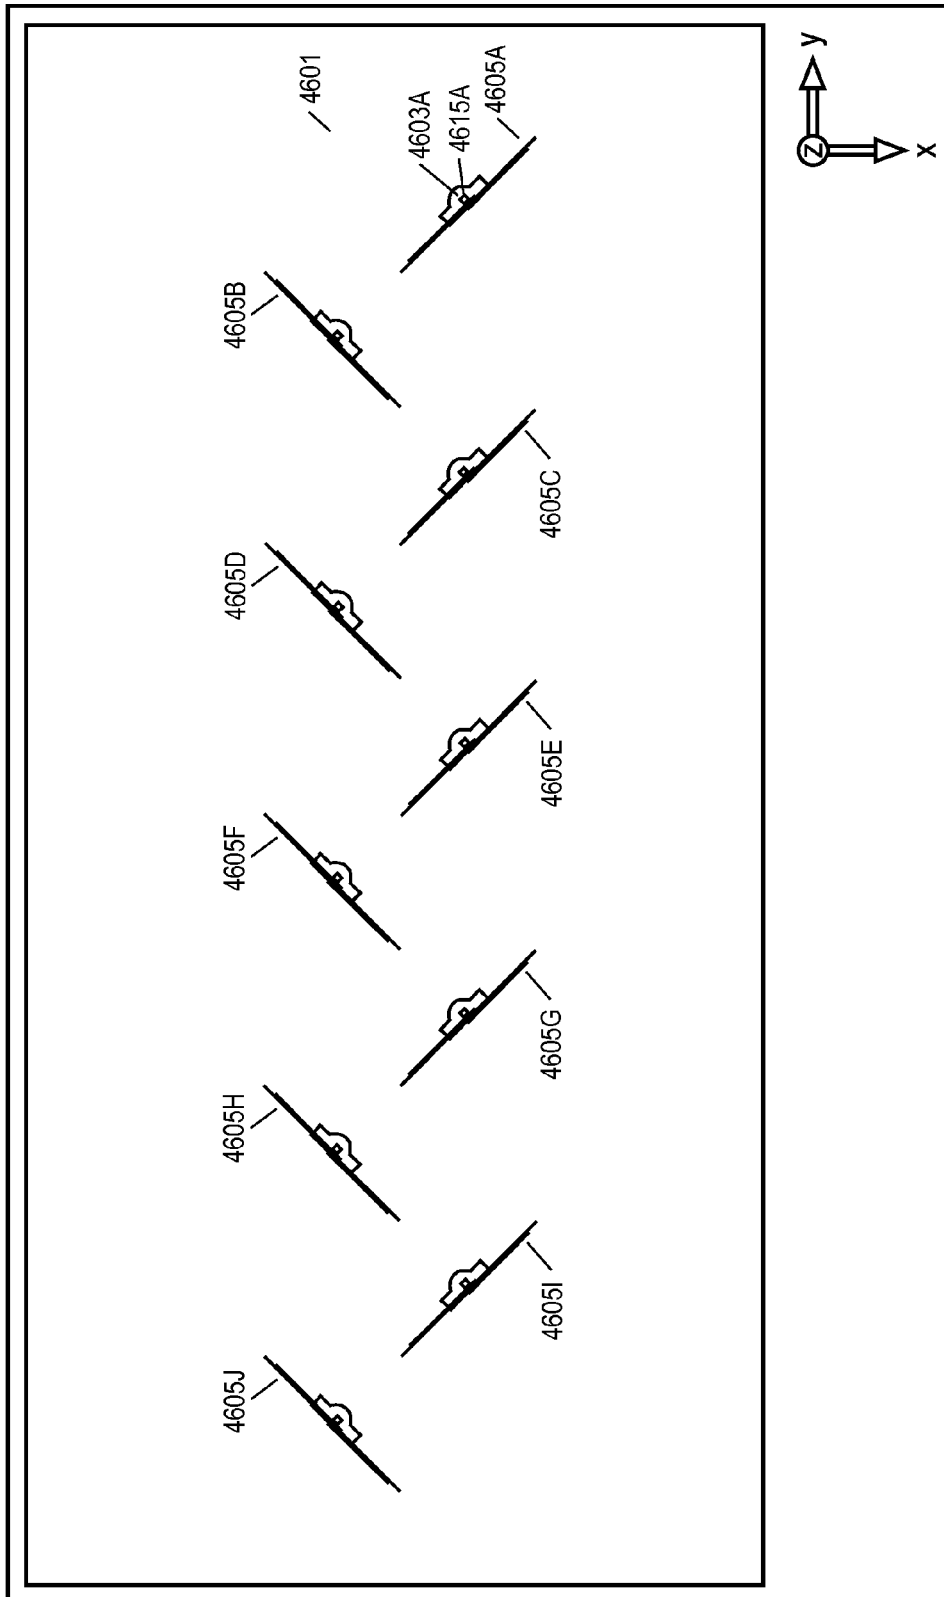


Fig. 46B




Fig. 46C




Fig. 47A




Fig. 47B




Fig. 48




Fig. 49A




Fig. 49B

5000




Fig. 50




Fig. 51




Fig. 52A




FIG. 52B




Fig. 52C




Fig. 53A




Fig. 53B




Fig. 53C




Fig. 53D

BACKHAUL RADIO WITH HORIZONTALLY OR VERTICALLY ARRANGED RECEIVE ANTENNA ARRAYS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a Continuation Application of U.S. application Ser. No. 14/336,958, filed on Jul. 21, 2014, which is a Continuation Application of U.S. application Ser. No. 13/898,429, filed on May 20, 2013 (U.S. Pat. No. 8,824,442), which is a Continuation of U.S. application Ser. No. 13/536,927, filed on Jun. 28, 2012 (U.S. Pat. No. 8,467,363), which is a Continuation-in-Part Application of U.S. application Ser. No. 13/371,366, filed on Feb. 10, 2012 (U.S. Pat. No. 8,311,023), which is a Continuation Application of U.S. application Ser. No. 13/212,036, filed on Aug. 17, 2011, (U.S. Pat. No. 8,238,318), and the disclosures of which are hereby incorporated herein by reference in their entireties.

BACKGROUND

1. Field

The present disclosure relates generally to data networking and in particular to a backhaul radio for connecting remote edge access networks to core networks and an associated antenna system.

2. Related Art

Data networking traffic has grown at approximately 100% per year for over 20 years and continues to grow at this pace. Only transport over optical fiber has shown the ability to keep pace with this ever-increasing data networking demand for core data networks. While deployment of optical fiber to an edge of the core data network would be advantageous from a network performance perspective, it is often impractical to connect all high bandwidth data networking points with optical fiber at all times. Instead, connections to remote edge access networks from core networks are often achieved with wireless radio, wireless infrared, and/or copper wireline technologies.

Radio, especially in the form of cellular or wireless local area network (WLAN) technologies, is particularly advantageous for supporting mobility of data networking devices. However, cellular base stations or WLAN access points inevitably become very high data bandwidth demand points that require continuous connectivity to an optical fiber core network.

When data aggregation points, such as cellular base station sites, WLAN access points, or other local area network (LAN) gateways, cannot be directly connected to a core optical fiber network, then an alternative connection, using, for example, wireless radio or copper wireline technologies, must be used. Such connections are commonly referred to as "backhaul."

Many cellular base stations deployed to date have used copper wireline backhaul technologies such as T1, E1, DSL, etc. when optical fiber is not available at a given site. However, the recent generations of HSPA+ and LTE cellular base stations have backhaul requirements of 100 Mb/s or more, especially when multiple sectors and/or multiple mobile network operators per cell site are considered. WLAN access points commonly have similar data backhaul requirements. These backhaul requirements cannot be practically satisfied at ranges of 300 m or more by existing copper wireline technologies. Even if LAN technologies such as Ethernet over multiple dedicated twisted pair wiring or hybrid fiber/coax technologies such as cable modems are considered, it is

impractical to backhaul at such data rates at these ranges (or at least without adding intermediate repeater equipment). Moreover, to the extent that such special wiring (i.e., CAT 5/6 or coax) is not presently available at a remote edge access network location; a new high capacity optical fiber is advantageously installed instead of a new copper connection.

Rather than incur the large initial expense and time delay associated with bringing optical fiber to every new location, it has been common to backhaul cell sites, WLAN hotspots, or LAN gateways from offices, campuses, etc. using microwave radios. An exemplary backhaul connection using the microwave radios **132** is shown in FIG. **1**. Traditionally, such microwave radios **132** for backhaul have been mounted on high towers **112** (or high rooftops of multi-story buildings) as shown in FIG. **1**, such that each microwave radio **132** has an unobstructed line of sight (LOS) **136** to the other. These microwave radios **132** can have data rates of 100 Mb/s or higher at unobstructed LOS ranges of 300 m or longer with latencies of 5 ms or less (to minimize overall network latency).

Traditional microwave backhaul radios **132** operate in a Point to Point (PTP) configuration using a single "high gain" (typically >30 dBi or even >40 dBi) antenna at each end of the link **136**, such as, for example, antennas constructed using a parabolic dish. Such high gain antennas mitigate the effects of unwanted multipath self-interference or unwanted co-channel interference from other radio systems such that high data rates, long range and low latency can be achieved. These high gain antennas however have narrow radiation patterns.

Furthermore, high gain antennas in traditional microwave backhaul radios **132** require very precise, and usually manual, physical alignment of their narrow radiation patterns in order to achieve such high performance results. Such alignment is almost impossible to maintain over extended periods of time unless the two radios have a clear unobstructed line of sight (LOS) between them over the entire range of separation. Furthermore, such precise alignment makes it impractical for any one such microwave backhaul radio to communicate effectively with multiple other radios simultaneously (i.e., a "point to multipoint" (PMP) configuration).

In wireless edge access applications, such as cellular or WLAN, advanced protocols, modulation, encoding and spatial processing across multiple radio antennas have enabled increased data rates and ranges for numerous simultaneous users compared to analogous systems deployed 5 or 10 years ago for obstructed LOS propagation environments where multipath and co-channel interference were present. In such systems, "low gain" (usually <6 dBi) antennas are generally used at one or both ends of the radio link both to advantageously exploit multipath signals in the obstructed LOS environment and allow operation in different physical orientations as would be encountered with mobile devices. Although impressive performance results have been achieved for edge access, such results are generally inadequate for emerging backhaul requirements of data rates of 100 Mb/s or higher, ranges of 300 m or longer in obstructed LOS conditions, and latencies of 5 ms or less.

In particular, "street level" deployment of cellular base stations, WLAN access points or LAN gateways (e.g., deployment at street lamps, traffic lights, sides or rooftops of single or low-multiple story buildings) suffers from problems because there are significant obstructions for LOS in urban environments (e.g., tall buildings, or any environments where tall trees or uneven topography are present).

FIG. **1** illustrates edge access using conventional unobstructed LOS PTP microwave radios **132**. The scenario depicted in FIG. **1** is common for many 2nd Generation (2G)

and 3rd Generation (3G) cellular network deployments using “macrocells”. In FIG. 1, a Cellular Base Transceiver Station (BTS) **104** is shown housed within a small building **108** adjacent to a large tower **112**. The cellular antennas **116** that communicate with various cellular subscriber devices **120** are mounted on the towers **112**. The PTP microwave radios **132** are mounted on the towers **112** and are connected to the BTSs **104** via an nT1 interface. As shown in FIG. 1 by line **136**, the radios **132** require unobstructed LOS.

The BTS on the right **104a** has either an nT1 copper interface or an optical fiber interface **124** to connect the BTS **104a** to the Base Station Controller (BSC) **128**. The BSC **128** either is part of or communicates with the core network of the cellular network operator. The BTS on the left **104b** is identical to the BTS on the right **104a** in FIG. 1 except that the BTS on the left **104b** has no local wireline nT1 (or optical fiber equivalent) so the nT1 interface is instead connected to a conventional PTP microwave radio **132** with unobstructed LOS to the tower on the right **112a**. The nT1 interfaces for both BTSs **104a**, **104b** can then be backhauled to the BSC **128** as shown in FIG. 1.

FIG. 2 is a block diagram of the major subsystems of a conventional PTP microwave radio **200** for the case of Time-Division Duplex (TDD) operation, and FIG. 3 is a block diagram of the major subsystems of a conventional PTP microwave radio **300** for the case of Frequency-Division Duplex (FDD) operation.

As shown in FIG. 2 and FIG. 3, the conventional PTP microwave radio traditionally uses one or more (i.e. up to “n”) T1 interfaces **204** (or in Europe, E1 interfaces). These interfaces **204** are common in remote access systems such as 2G cellular base stations or enterprise voice and/or data switches or edge routers. The T1 interfaces are typically multiplexed and buffered in a bridge (e.g., the Interface Bridge **208**, **308**) that interfaces with a Media Access Controller (MAC) **212**, **312**.

The MAC **212**, **312** is generally denoted as such in reference to a sub-layer of Layer 2 within the Open Systems Interconnect (OSI) reference model. Major functions performed by the MAC include the framing, scheduling, prioritizing (or “classifying”), encrypting and error checking of data sent from one such radio at FIG. 2 or FIG. 3 to another such radio. The data sent from one radio to another is generally in a “user plane” if it originates at the T1 interface(s) or in the “control plane” if it originates internally such as from the Radio Link Controller (RLC) **248**, **348** shown in FIG. 2 or FIG. 3. A typical MAC frame format **400** (known as a MAC protocol data unit, or “MPDU”) with header **404**, frame body **408** and frame check sum (FCS) **412** is shown in FIG. 4.

With reference to FIGS. 2 and 3, the Modem **216**, **316** typically resides within the “baseband” portion of the Physical (PHY) layer 1 of the OSI reference model. In conventional PTP radios, the baseband PHY, depicted by Modem **216**, **316**, typically implements scrambling, forward error correction encoding, and modulation mapping for a single RF carrier in the transmit path. In receive, the modem typically performs the inverse operations of demodulation mapping, decoding and descrambling. The modulation mapping is conventionally Quadrature Amplitude Modulation (QAM) implemented with In-phase (I) and Quadrature-phase (Q) branches.

The Radio Frequency (RF) **220**, **320** also resides within the PHY layer of the radio. In conventional PTP radios, the RF **220**, **320** typically includes a single transmit chain (Tx) **224**, **324** that includes I and Q digital to analog converters (DACs), a vector modulator, optional upconverters, a programmable gain amplifier, one or more channel filters, and one or more combinations of a local oscillator (LO) and a frequency syn-

thesizer. Similarly, the RF **220**, **320** also typically includes a single receive chain (Rx) **228**, **328** that includes I and Q analog to digital converters (ADCs), one or more combinations of an LO and a frequency synthesizer, one or more channel filters, optional downconverters, a vector demodulator and an automatic gain control (AGC) amplifier. Note that in many cases some of the one or more LO and frequency synthesizer combinations can be shared between the Tx and Rx chains.

As shown in FIGS. 2 and 3, conventional PTP radios **200**, **300** also include a single power amplifier (PA) **232**, **332**. The PA **232**, **332** boosts the transmit signal to a level appropriate for radiation from the antenna in keeping with relevant regulatory restrictions and instantaneous link conditions. Similarly, such conventional PTP radios **232**, **332** typically also include a single low-noise amplifier (LNA) **236**, **336** as shown in FIGS. 2 and 3. The LNA **236**, **336** boosts the received signal at the antenna while minimizing the effects of noise generated within the entire signal path.

As described above, FIG. 2 illustrates a conventional PTP radio **200** for the case of TDD operation. As shown in FIG. 2, conventional PTP radios **200** typically connect the antenna **240** to the PA **232** and LNA **236** via a band-select filter **244** and a single-pole, single-throw (SPST) switch **242**.

As described above, FIG. 3 illustrates a conventional PTP radio **300** for the case of FDD operation. As shown in FIG. 3, in conventional PTP radios **300**, then antenna **340** is typically connected to the PA **332** and LNA **336** via a duplexer filter **344**. The duplexer filter **344** is essentially two band-select filters (tuned respectively to the Tx and Rx bands) connected at a common point.

In the conventional PTP radios shown in FIGS. 2 and 3, the antenna **240**, **340** is typically of very high gain such as can be achieved by a parabolic dish so that gains of typically >30 dBi (or even sometimes >40 dBi), can be realized. Such an antenna usually has a narrow radiation pattern in both the elevation and azimuth directions. The use of such a highly directive antenna in a conventional PTP radio link with unobstructed LOS propagation conditions ensures that the modem **216**, **316** has insignificant impairments at the receiver (antenna **240**, **340**) due to multipath self-interference and further substantially reduces the likelihood of unwanted co-channel interference due to other nearby radio links.

Although not explicitly shown in FIGS. 2 and 3, the conventional PTP radio may use a single antenna structure with dual antenna feeds arranged such that the two electromagnetic radiation patterns emanated by such an antenna are nominally orthogonal to each other. An example of this arrangement is a parabolic dish. Such an arrangement is usually called dual-polarized and can be achieved either by orthogonal vertical and horizontal polarizations or orthogonal left-hand circular and right-hand circular polarizations.

When duplicate modem blocks, RF blocks, and PA/LNA/switch blocks are provided in a conventional PTP radio, then connecting each PHY chain to a respective polarization feed of the antenna allows theoretically up to twice the total amount of information to be communicated within a given channel bandwidth to the extent that cross-polarization self-interference can be minimized or cancelled sufficiently. Such a system is said to employ “dual-polarization” signaling.

When an additional circuit (not shown) is added to FIG. 2 that can provide either the RF Tx signal or its anti-phase equivalent to either one or both of the two polarization feeds of such an antenna, then “cross-polarization” signaling can be used to effectively expand the constellation of the modem within any given symbol rate or channel bandwidth. With two polarizations and the choice of RF signal or its anti-phase,

5

then an additional two information bits per symbol can be communicated across the link. Theoretically, this can be extended and expanded to additional phases, representing additional information bits. At the receiver, for example, a circuit (not shown) could detect if the two received polarizations are anti-phase with respect to each other, or not, and then combine appropriately such that the demodulator in the modem block can determine the absolute phase and hence deduce the values of the two additional information bits. Cross-polarization signaling has the advantage over dual-polarization signaling in that it is generally less sensitive to cross-polarization self-interference but for high order constellations such as 64-QAM or 256-QAM, the relative increase in channel efficiency is smaller.

In the conventional PTP radios shown in FIGS. 2 and 3, substantially all the components are in use at all times when the radio link is operative. However, many of these components have programmable parameters that can be controlled dynamically during link operation to optimize throughput and reliability for a given set of potentially changing operating conditions. The conventional PTP radios of FIGS. 2 and 3 control these link parameters via a Radio Link Controller (RLC) 248, 348. The RLC functionality is also often described as a Link Adaptation Layer that is typically implemented as a software routine executed on a microcontroller within the radio that can access the MAC 212, 312, Modem 216, 316, RF 220, 320 and/or possibly other components with controllable parameters. The RLC 248, 348 typically can both vary parameters locally within its radio and communicate with a peer RLC at the other end of the conventional PTP radio link via "control frames" sent by the MAC 212, 312 with an appropriate identifying field within a MAC Header 404 (in reference to FIG. 4).

Typical parameters controllable by the RLC 248, 348 for the Modem 216, 316 of a conventional PTP radio include encoder type, encoding rate, constellation selection and reference symbol scheduling and proportion of any given PHY Protocol Data Unit (PPDU). Typical parameters controllable by the RLC 248, 348 for the RF 220, 320 of a conventional PTP radio include channel frequency, channel bandwidth, and output power level. To the extent that a conventional PTP radio employs two polarization feeds within its single antenna, additional parameters may also be controlled by the RLC 248, 348 as self-evident from the description above.

In conventional PTP radios, the RLC 248, 348 decides, usually autonomously, to attempt such parameter changes for the link in response to changing propagation environment characteristics such as, for example, humidity, rain, snow, or co-channel interference. There are several well-known methods for determining that changes in the propagation environment have occurred such as monitoring the receive signal strength indicator (RSSI), the number of or relative rate of FCS failures at the MAC 212, 312, and/or the relative value of certain decoder accuracy metrics. When the RLC 248, 348 determines that parameter changes should be attempted, it is necessary in most cases that any changes at the transmitter end of the link become known to the receiver end of the link in advance of any such changes. For conventional PTP radios, and similarly for many other radios, there are at least two well-known techniques which in practice may not be mutually exclusive. First, the RLC 248, 348 may direct the PHY, usually in the Modem 216, 316 relative to FIGS. 2 and 3, to pre-pend a PHY layer convergence protocol (PLCP) header to a given PPDU that includes one or more (or a fragment thereof) given MPDUs wherein such PLCP header has information fields that notify the receiving end of the link of parameters used at the transmitting end of the link. Second,

6

the RLC 248, 348 may direct the MAC 212, 312 to send a control frame, usually to a peer RLC 248, 348, including various information fields that denote the link adaptation parameters either to be deployed or to be requested or considered.

The foregoing describes at an overview level the typical structural and operational features of conventional PTP radios which have been deployed in real-world conditions for many radio links where unobstructed (or substantially unobstructed) LOS propagation was possible. The conventional PTP radio on a whole is completely unsuitable for obstructed LOS or PMP operation.

SUMMARY

The following summary of the invention is included in order to provide a basic understanding of some aspects and features of the invention. This summary is not an extensive overview of the invention and as such it is not intended to particularly identify key or critical elements of the invention or to delineate the scope of the invention. Its sole purpose is to present some concepts of the invention in a simplified form as a prelude to the more detailed description that is presented below.

Some embodiments of the claimed invention are directed to backhaul radios that are compact, light and low power for street level mounting, operate at 100 Mb/s or higher at ranges of 300 m or longer in obstructed LOS conditions with low latencies of 5 ms or less, can support PTP and PMP topologies, use radio spectrum resources efficiently and do not require precise physical antenna alignment. Radios with such exemplary capabilities as described by multiple various embodiments are referred to herein by the term "Intelligent Backhaul Radio" (IBR). Exemplary IBRs are disclosed in detail in co-pending U.S. patent application Ser. No. 13/212,036, entitled Intelligent Backhaul Radio, filed Aug. 17, 2011, and U.S. patent application Ser. No. 13/271,051, entitled Intelligent Backhaul System, filed Oct. 11, 2011, U.S. patent application Ser. No. 13/415,778, entitled Intelligent Backhaul System, filed Mar. 8, 2012, U.S. patent application Ser. No. 13/448,294, entitled Hybrid Band Intelligent Backhaul Radio, filed Apr. 16, 2012, U.S. patent application Ser. No. 13/371,366, entitled Intelligent Backhaul Radio, filed Feb. 10, 2012, the entireties of which is hereby incorporated by reference.

The IBR provides reliable, high-capacity data communications in urban deployments. In some deployments, there is not line-of-sight visibility between the two end nodes of the link, and in such cases the propagation channel relies on multipath scattering and diffraction to achieve communication. The resulting wireless channel features high path loss exponents, time-domain delay spread, and wide angular-spread of incoming power (i.e. spatial Rayleigh fading). In addition, some embodiments of the IBR rely on two-stream spatial multiplexing, which places further demands on the antenna design.

To achieve the desired throughput and reliability in the above conditions and utilizing specific embodiments, the antenna system must meet the following specifications:

- Wide azimuthal coverage (e.g. 100 to 120 degrees) to accommodate for the wide-angular spread of incoming power;
- Narrow vertical beamwidths (e.g. 15 degrees);
- At minimum two uncorrelated antenna ports for any direction within the coverage zone;
- High-gain to accommodate the high-path loss; and
- Compact overall enclosure.

Fundamentally, there are three diversity mechanisms available to achieve uncorrelated signals at the terminals of an antenna: spatial diversity, polarization diversity, and angular beam diversity. The optimal selection or combination of these diversity mechanisms depends on a priori knowledge of typical channel characteristics.

For example, in a wide angle-of-arrival environment, uncorrelated signals can be achieved with omni-directional antennas with relatively compact spatial separation (on the order of one wavelength) e.g. most consumer-grade WiFi access points utilize spatial diversity, as known in the art. This spacing requirement increases as the channel becomes more directional, making spatial diversity a poor choice for directional scenarios when minimal physical size is important.

Polarization diversity has the advantage that a well-designed dual-polarity antenna can be achieved in virtually the same overall dimensions as the equivalent single-polarization antenna. Polarization diversity is advantageous where there is a channel with significant polarization-dependent scattering. Polarization diversity typically requires that the polarizations be orthogonal, i.e. the relative angle between the two antennas should be 90 degrees; however, the two polarizations can be arbitrarily rotated with regards to the earth.

Angular beam diversity is particularly effective if the incoming angular power distribution is clustered around more than one dominant angles of arrival. The angle between these clusters dictates the minimum angular width of the antenna pattern that will achieve uncorrelated antennas. Angular beam diversity has two advantages over spatial diversity, gain and compactness. As the beamwidth of an antenna pattern narrows the gain of the antenna increases, and the sources for each beam can be co-located without any loss in diversity. For example, in analog beamforming schemes such as a butler matrix, the same antenna elements are re-used to form the various directional beams.

According to an aspect of the invention, an intelligent backhaul radio is disclosed that includes a directional transmit antenna array, for simultaneously transmitting a plurality of transmit symbol streams to a target radio, said directional transmit antenna array comprising a plurality of transmit antenna sub-arrays, wherein each of said plurality of transmit antenna sub-arrays comprise a plurality of transmit antenna elements; a plurality of directional receive antenna arrays, for simultaneously receiving a plurality of receive RF signals comprising a plurality of receive symbol streams from the target radio, wherein each of said plurality of directional receive antenna arrays comprise a plurality of receive antenna sub-arrays, wherein each of said plurality of receive antenna sub-arrays comprise a plurality of receive antenna elements, wherein each of said receive antenna sub-arrays has a normalized antenna pattern similar to the normalized antenna pattern of each of the transmit antenna sub-arrays in elevation and narrower than the normalized antenna pattern of each of the transmit antenna sub-arrays in azimuth, and wherein each of the plurality of directional receive antenna arrays are positioned physically offset from said directional transmit antenna array, and distributed in azimuthal orientation such that the aggregate normalized azimuthal antenna patterns of the plurality of receive antenna sub-arrays in composite is equal to or greater than the normalized antenna pattern of each of the transmit antenna sub-arrays in azimuth; one or more demodulator cores, wherein each demodulator core demodulates one or more of said plurality of receive symbol streams to produce a respective receive data interface stream; a plurality of receive RF chains to convert the plurality of receive RF signals to a plurality of respective receive chain output signals; a frequency selective receive path channel

multiplexer between the one or more demodulator cores and the plurality of receive RF chains, the frequency selective receive path channel multiplexer to generate the one or more receive symbol streams from the plurality of receive chain output signals; and one or more selectable RF connections that selectively couple certain of the plurality of said receive antenna sub-arrays to certain of the plurality of receive RF chains, wherein the number of receive antenna sub-arrays that can be selectively coupled to receive RF chains exceeds the number of receive RF chains that can accept receive RF signals from the one or more selectable RF connections; and a radio resource controller, wherein the radio resource controller sets or causes to be set the specific selective couplings between the certain of the plurality of receive antenna sub-arrays and the certain of the plurality of receive RF chains.

Each of said plurality of transmit sub-arrays may transmit one or more of said plurality of transmit symbol streams, and at least two of said plurality of transmit sub-arrays may utilize differing polarizations.

The plurality of transmit antenna elements may be printed dipole antenna elements and said plurality of receive antenna elements may be patch antenna elements.

One or more of said receive patch antenna elements may be utilized by a plurality of the receive antenna sub-arrays of the directional receive antenna array, said receive patch antenna element utilizing separate antenna feed structures providing for substantially orthogonally polarized reception for each of the sub-arrays.

At least one of said one or more of said plurality of transmit symbol streams may be multiplexed into a plurality of transmit RF signals. The plurality of transmit RF signals may be respectively transmitted from each of said plurality of transmit antenna sub-arrays. At least two of said RF signals may include the same transmit symbol stream.

The same transmit symbol stream may be modified in gain, phase, or delay prior to transmission by each of said plurality of transmit sub arrays to achieve a composite transmit symbol stream property.

The composite transmit symbol stream property may be a modified polarization.

The composite transmit symbol stream property may result in an improved receive property at a receiver of the target radio.

The improved receive property may be one or more of: a received RSSI, a received signal to noise ratio, a signal to interference ratio, a de-correlation between differing transmit streams, a propagation channel property, a delay spread, and a signal variability.

A metric may be provided by said target radio to the intelligent backhaul radio, said metric used to adapt said gain, phase, or delay of said transmit symbol stream prior to transmission by each of said plurality of transmit sub arrays, wherein said metric is derived from the receive property at the target radio.

At least one of said plurality of transmit symbol streams may include a unique transmit symbol stream.

The plurality of transmit antenna sub-arrays of said directional transmit antenna array may utilize a common ground plane, and said ground plane may include non-conducting slots at the edge of said ground plane. The slots may be aligned with said printed dipole antenna elements of a first of the transmit sub-arrays to effect an associated antenna pattern of that sub-array, and substantially not affect the antenna pattern of a second transmit sub-array of said plurality of transmit sub-arrays utilizing said shared ground plane. The

second transmit sub-array may include antenna elements having a polarization orthogonal to the polarization of said first transmit sub-array.

The plurality of receive antenna arrays positioned physically offset from said directional transmit antenna array may be distributed linearly. The linear distribution of said plurality of directional receive antenna arrays may be vertical. The linear distribution of said plurality of directional receive antenna arrays may be horizontal.

The plurality of transmit antenna sub-arrays of said directional transmit antenna array may utilize a common ground plane and each of said printed dipole antenna elements of a first transmit sub-array of said plurality of transmit sub-arrays may utilize a first polarization of said at least two differing polarizations, and may be printed on a single printed circuit board.

Each of said printed dipole antenna elements of a second sub-array of said plurality of transmit sub-arrays may utilize a second polarization of said at least two differing polarizations, and may be printed on individual printed circuit boards.

The single printed circuit board of said first sub-array may include slots for mechanically aligning, securing, or joining with complementary slots located within said individual printed circuit boards of said second sub-array.

Each of the individual printed circuit boards and the single printed circuit board may include tabs for mating with slots in a printed circuit board of said common ground plane and for mechanically securing said printed dipole antenna elements.

Each of the individual printed circuit boards and the single printed circuit board may include tabs for mating with slots in a printed circuit board of said common ground plane and for coupling feed networks disposed upon said printed circuit board of said common ground plane with said printed dipole antenna elements.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more examples of embodiments and, together with the description of example embodiments, serve to explain the principles and implementations of the embodiments.

FIG. 1 is an illustration of conventional point to point (PTP) radios deployed for cellular base station backhaul with unobstructed line of sight (LOS).

FIG. 2 is a block diagram of a conventional PTP radio for Time Division Duplex (TDD).

FIG. 3 is a block diagram of a conventional PTP radio for Frequency Division Duplex (FDD).

FIG. 4 is an illustration of a MAC Protocol Data Unit (MPDU).

FIG. 5 is an illustration of intelligent backhaul radios (IBRs) deployed for cellular base station backhaul with obstructed LOS according to one embodiment of the invention.

FIG. 6 is a block diagram of an IBR according to one embodiment of the invention.

FIG. 7 is a block diagram of an IBR according to one embodiment of the invention.

FIG. 8 is a block diagram illustrating an exemplary deployment of IBRs according to one embodiment of the invention.

FIG. 9 is a block diagram illustrating an exemplary deployment of IBRs according to one embodiment of the invention.

FIG. 10 is a block diagram of an IBR antenna array according to one embodiment of the invention.

FIG. 10A is a block diagram of an IBR antenna array according to one embodiment of the invention.

FIG. 11 is a block diagram of a front-end unit for TDD operation according to one embodiment of the invention.

FIG. 12 is a block diagram of a front-end unit for FDD operation according to one embodiment of the invention.

FIG. 12A is a block diagram of a front-end transmission unit according to one embodiment of the invention.

FIG. 12B is a block diagram of a front-end reception unit according to one embodiment of the invention.

FIG. 13 is a perspective view of an IBR according to one embodiment of the invention.

FIG. 14 is a perspective view of an IBR according to one embodiment of the invention.

FIG. 15 is a perspective view of an IBR according to one embodiment of the invention.

FIG. 16 is a block diagram illustrating an exemplary transmit chain within an IBR RF according to one embodiment of the invention.

FIG. 17 is a block diagram illustrating an exemplary receive chain within an IBR RF according to one embodiment of the invention.

FIG. 18 is a block diagram illustrating an IBR modem according to one embodiment of the invention.

FIG. 19 is a block diagram illustrating a modulator core j according to one embodiment of the invention.

FIG. 20 is a block diagram illustrating a demodulator core j according to one embodiment of the invention.

FIG. 21 is a block diagram illustrating a modulator core j according to one embodiment of the invention.

FIG. 22 is a block diagram illustrating a demodulator core j according to one embodiment of the invention.

FIG. 23, consisting of FIG. 23A and FIG. 23B, is a block diagram illustrating a channel multiplexer (MUX) according to one embodiment of the invention. FIG. 23A is a partial view showing the transmit path and the channel equalizer coefficients generator within the exemplary channel MUX. FIG. 23B is a partial view showing the receive path within the exemplary channel MUX.

FIG. 24 is a block diagram illustrating an exemplary Tx-CE-m according to one embodiment of the invention.

FIG. 24A is a block diagram illustrating an exemplary Tx-CE-m, including a complex circular FIR capability, according to one embodiment of the invention.

FIG. 24B is a plot of an exemplary desired transmit equalizer impulse response for use with the Tx-CE-m of FIG. 24A according to one embodiment of the invention.

FIG. 24C is a plot of an exemplary desired transmit equalizer frequency response for use with the Tx-CE-m of FIG. 24A according to one embodiment of the invention.

FIG. 24D is a plot of an exemplary desired transmit equalizer impulse response for use with the Tx-CE-m of FIG. 24A according to one embodiment of the invention.

FIG. 24E is a plot of an exemplary 16 QAM-10 symbol Transmit Block ($\overrightarrow{\text{TxBlock}}_k$) for processing with the Tx-CE-m of FIG. 24A according to one embodiment of the invention.

FIG. 24F is a plot of the result of a linear convolution of the exemplary 16 QAM-10 symbol Transmit Block ($\overrightarrow{\text{TxBlock}}_k$) according to one embodiment of the invention.

FIG. 24G is a plot of the result of the FIR based circular convolution according to one embodiment of the invention.

FIG. 24H is a plot of the result of a frequency domain based circular convolution of the exemplary 16 QAM-10 symbol Transmit Block ($\overrightarrow{\text{TxBlock}}_k$) according to one embodiment of the invention.

11

FIG. 24I is a magnitude plot of the result of both the circular FIR and the frequency domain based circular convolutions of the exemplary 16 QAM-10 symbol Transmit Block ($\overrightarrow{TxBlock_k}$) according to one embodiment of the invention.

FIG. 25 is a block diagram illustrating an exemplary Rx-CE-I according to one embodiment of the invention.

FIG. 26 is a timing diagram illustrating processing of PPDU-I with Tx-path and Rx-path of respective IBR channel MUXs according to one embodiment of the invention.

FIG. 27 is a block diagram illustrating an exemplary Tx PLCP according to one embodiment of the invention.

FIG. 28 is a block diagram illustrating an exemplary Tx PLCP Mod-j according to one embodiment of the invention.

FIG. 29 is a block diagram illustrating an exemplary Rx PLCP according to one embodiment of the invention.

FIG. 30 is a block diagram illustrating an exemplary Rx PLCP Mod-j according to one embodiment of the invention.

FIG. 31 is a schematic diagram of an IBR communications protocols stack according to one embodiment of the invention.

FIG. 32 is a schematic diagram of an IBR communications protocols stack according to one embodiment of the invention.

FIG. 33 is a block diagram of an IBR media access control (MAC) according to one embodiment of the invention.

FIG. 34 is a timing diagram illustrating channel activity for FDD with fixed superframe timing according to one embodiment of the invention.

FIG. 35 is a timing diagram illustrating channel activity for TDD with fixed superframe timing according to one embodiment of the invention.

FIG. 36 is a timing diagram illustrating channel activity for TDD/CSMA with variable superframe timing according to one embodiment of the invention.

FIG. 37 is a diagram illustrating a sector antenna panel and processing according to one embodiment of the invention.

FIG. 38 is a plot of an exemplary far field antenna pattern of a sector antenna panel according to one embodiment of the invention.

FIG. 39 is a diagram illustrating an alternative sector antenna panel and processing, including separate transmit and receive antennas according to one embodiment of the invention.

FIG. 40 is a plot of an exemplary far field antenna pattern of an alternative sector antenna panel including separate transmit and receive antennas according to one embodiment of the invention.

FIG. 41 is a diagram illustrating an alternative sector antenna panel and processing including separate transmit and receive antennas according to one embodiment of the invention.

FIG. 42 is a plot of an exemplary far field antenna pattern of an alternative sector antenna panel and processing including separate transmit and receive antennas according to one embodiment of the invention.

FIG. 43 is a diagram of a dual-polarity, two-port patch antenna element including feed and grounding points according to one embodiment of the invention.

FIG. 44A is a diagram of a front view of an exemplary dual-polarity, two port, patch antenna array according to one embodiment of the invention.

FIG. 44B is a diagram of a side view of an exemplary dual-polarity, two port, patch antenna array according to one embodiment of the invention.

12

FIG. 45A is a diagram of a view of an exemplary single-polarity, single port, printed dipole antenna element for use in an antenna array according to one embodiment of the invention.

FIG. 45B is a diagram of an alternative view of an exemplary single-polarity, single port, printed dipole antenna element for use in an antenna array according to one embodiment of the invention.

FIG. 46A is a diagram of a view of an exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements according to one embodiment of the invention.

FIG. 46B is a diagram of a front view of an exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements according to one embodiment of the invention.

FIG. 46C is a diagram of a back view of an exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements, showing a cooperate feed network for each single polarized sub-array associated with each port according to one embodiment of the invention.

FIG. 47A is a plot of an exemplary far field elevation antenna pattern of a dual-polarity, two port, antenna array utilizing printed dipole antenna elements according to one embodiment of the invention.

FIG. 47B is a plot of an exemplary far field azimuthal antenna pattern of a dual-polarity, two port, antenna array utilizing printed dipole antenna elements according to one embodiment of the invention.

FIG. 48 is a diagram of a view of an alternative exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements, including alternative printed dipole antenna elements and further providing vertical and horizontal polarizations according to one embodiment of the invention.

FIG. 49A is a diagram of an alternative exemplary printed dipole antenna element according to one embodiment of the invention.

FIG. 49B is a diagram of an alternative exemplary printed dipole antenna including illustrative ground plane interface according to one embodiment of the invention.

FIG. 50 is a diagram of an exemplary printed dipole antenna structure for use in an antenna array according to one embodiment of the invention.

FIG. 51 is a diagram depicting an exemplary assembly a dual-polarized, two port, printed dipole antenna array utilizing a printed dipole structure according to one embodiment of the invention.

FIG. 52A is a diagram of a front view of an exemplary horizontally arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

FIG. 52B is a diagram of an alternative view of an exemplary horizontally arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

FIG. 52C is a diagram of a top of an exemplary horizontally arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

FIG. 53A is a diagram of a front view of an exemplary vertically arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

FIG. 53B is a diagram of an alternative view of an exemplary vertically arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

FIG. 53C is a diagram of a side view of an exemplary vertically arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

13

FIG. 53D is a diagram of a top view of an exemplary vertically arranged intelligent backhaul radio antenna array according to one embodiment of the invention.

DETAILED DESCRIPTION

FIG. 5 illustrates deployment of intelligent backhaul radios (IBRs) in accordance with an embodiment of the invention. As shown in FIG. 5, the IBRs 500 are deployable at street level with obstructions such as trees 504, hills 508, buildings 512, etc. between them. The IBRs 500 are also deployable in configurations that include point to multipoint (PMP), as shown in FIG. 5, as well as point to point (PTP). In other words, each IBR 500 may communicate with more than one other IBR 500.

For 3G and especially for 4th Generation (4G), cellular network infrastructure is more commonly deployed using “microcells” or “picocells.” In this cellular network infrastructure, compact base stations (eNodeBs) 516 are situated outdoors at street level. When such eNodeBs 516 are unable to connect locally to optical fiber or a copper wireline of sufficient data bandwidth, then a wireless connection to a fiber “point of presence” (POP) requires obstructed LOS capabilities, as described herein.

For example, as shown in FIG. 5, the IBRs 500 include an Aggregation End IBR (AE-IBR) and Remote End IBRs (RE-IBRs). The eNodeB 516 of the AE-IBR is typically connected locally to the core network via a fiber POP 520. The RE-IBRs and their associated eNodeBs 516 are typically not connected to the core network via a wireline connection; instead, the RE-IBRs are wirelessly connected to the core network via the AE-IBR. As shown in FIG. 5, the wireless connection between the IBRs include obstructions (i.e., there may be an obstructed LOS connection between the RE-IBRs and the AE-IBR).

FIGS. 6 and 7 illustrate exemplary embodiments of the IBRs 500 shown in FIG. 5. In FIGS. 6 and 7, the IBRs 500 include interfaces 604, interface bridge 608, MAC 612, modem 624, channel MUX 628, RF 632, which includes Tx1 . . . TxM 636 and Rx1 . . . RxN 640, antenna array 648 (includes multiple antennas 652), a Radio Link Controller (RLC) 656 and a Radio Resource Controller (RRC) 660. The IBR may optionally include an IBMS agent 700 as shown in FIG. 7. It will be appreciated that the components and elements of the IBRs may vary from that illustrated in FIGS. 6 and 7. The various elements of the IBR are described below by their structural and operational features in numerous different embodiments.

The external interfaces of the IBR (i.e., the IBR Interface Bridge 608 on the wireline side and the IBR Antenna Array 648 (including antennas 652) on the wireless side) are a starting point for describing some fundamental differences between the numerous different embodiments of the IBR 500 and the conventional PTP radios described above (or other commonly known radio systems, such as those built to existing standards including 802.11n (WiFi) and 802.16e (WiMax)).

In some embodiments, the IBR Interface Bridge 608 physically interfaces to standards-based wired data networking interfaces 604 as Ethernet 1 through Ethernet P. “P” represents a number of separate Ethernet interfaces over twisted-pair, coax or optical fiber. The IBR Interface Bridge 608 can multiplex and buffer the P Ethernet interfaces 604 with the IBR MAC 612. For the case of P=1 (a single Ethernet interface), the multiplexing operation of the IBR Interface Bridge 608 is a trivial “pass-through” between the single Ethernet interface and the buffer. In exemplary embodiments, the IBR

14

Interface Bridge 608 preserves “Quality of Service” (QoS) or “Class of Service” (CoS) prioritization as indicated, for example, in IEEE 802.1q 3-bit Priority Code Point (PCP) fields within the Ethernet frame headers, such that either the IBR MAC 612 schedules such frames for transmission according to policies configured within the IBR of FIG. 6 or communicated via the IBMS Agent 700 of FIG. 7, or the IBR interface bridge 608 schedules the transfer of such frames to the IBR MAC 612 such that the same net effect occurs. In other embodiments, the IBR interface bridge 608 also forwards and prioritizes the delivery of frames to or from another IBR over an instant radio link based on Multiprotocol Label Switching (MPLS) or Multiprotocol Label Switching Transport Profile (MPLS-TP).

In some embodiments, the IBR Interface Bridge 608 can also perform layer 2 switching of certain Ethernet interfaces to other Ethernet interfaces 604 in response to radio link failure conditions and policies configured within the IBR of FIG. 6 or communicated via the IBMS Agent 700 of FIG. 7. A radio link failure condition can arise from any one or more of multiple possible events, such as the following exemplary radio link failure condition events:

- physical failure of a component within the IBR other than the IBR Interface Bridge and its power supply;
- degradation of the RF link beyond some pre-determined throughput level due to either changing propagation environment or additional co-channel interference; and
- failure of any kind at the other end of the RF link that prevents connection to the ultimate source or destination.

FIG. 8 illustrates an exemplary configuration of multiple IBRs (IBR 1 804, IBR 2 808). Each IBR 804, 808 has layer 2 switching in the IBR interface bridges 816, 820 (each corresponding to an instance of the IBR Interface Bridge 608 of FIGS. 6 and 7). In one embodiment, the data networking traffic can be from, for example, a cellular eNodeB, a WiFi hotspot, a enterprise edge router, or any of numerous other remote data networks 828. As shown in FIG. 8, the remote data network 828 is connected via an Ethernet cable (copper or fiber) 832 to the first Ethernet port 836 on the IBR Interface Bridge 816 of IBR1 804. Another Ethernet cable 840 connects the second Ethernet port 844 on the IBR Interface Bridge 816 of IBR1 804 to the first Ethernet port 848 on the IBR Interface Bridge 820 of IBR2 808. If a radio link failure condition occurs for any reason, such as those listed above, with respect to RF Link 1 852, then the layer 2 switch within the IBR Interface Bridge 816 of IBR1 804 can automatically connect all data networking traffic originating from or destined to Remote Data Network 1 828 via IBR2 808 and RF Link2 856, completely transparently to Remote Data Network 1 828. This provides fail over redundancy to reduce the probability of network outage at Remote Data Network 1 828.

In some embodiments, the IBR Interface Bridge with layer 2 switching can also be configured to perform load balancing in response to operating conditions and policies configured within the IBR of FIG. 6 or communicated via the IBMS Agent 700 of FIG. 7. For example, with reference to FIG. 8, the layer 2 switch in the IBR Interface Bridge 816 of IBR1 804 can connect all data networking traffic in excess of data rates above a certain limit, such as a pre-determined level or the instantaneous supportable rate on RF Link 1 852, over to IBR2 808. For full two-way functionality of this load balancing, an analogous load balancing capability exists within the layer 2 switch of an IBR at the respective other ends of both RF Link 1 852 and RF Link 2 856.

FIG. 9 illustrates an alternative configuration of IBRs 804, 808 with layer 2 switching capability and optional load bal-

15

ancing capability for the case of two disparate Remote Data Networks 1 and 2 (900 and 904, respectively). The two Remote Data Networks 1 and 2 (900 and 904) are at or within Ethernet cabling distance of two IBRs 1 and 2 (804, 808) operating on two RF links 1 and 2 (852, 856). As described above, the Remote Data Network 828 is connected via Ethernet cable 832 to the first Ethernet port 836, and the IBRs 1 and 2 (804, 808) are connected via Ethernet cable 840 at Ethernet ports 844 and 848 respectively. The Remote Data Network 2 904 is connected to IBR 2 808 via Ethernet cable 908 at the Ethernet port 912. In the embodiment shown in FIG. 9, if a radio link failure condition occurs for any reason, such as those listed above, with respect to RF Link 1 852, then the IBR Interface Bridge 816 of IBR1 804 can use its layer 2 switch to connect to the IBR Interface Bridge 820 of IBR2 808 via Ethernet cable 840 such that IBR2 808 can backhaul, subject to its throughput capabilities and prioritization policies, the traffic originating from or destined to both Remote Data Network 1 900 and Remote Data Network 2 904. Similarly, the IBRs 804, 808 can perform load balancing across both RF Links 1 and 2 (852, 856), for traffic originating from or destined to Remote Data Networks 1 and 2 (900, 904).

In some embodiments, RF link 1 852 may utilize spectrum possessing advantageous conditions, such as reduced interference, wider channel bandwidth, better propagation characteristics, and/or higher allowable power than the spectrum utilized by RF Link 2 856, or vice versa. In the situation where a radio link failure condition occurs with respect to the more advantageous spectrum, either control signaling between the two IBR Interface Bridges 816, 820 of IBRs 1 and 2 as shown in FIG. 6 or messaging between the two IBMS Agents 700 as shown in FIG. 7, whether directly or indirectly via one or more intermediaries, can cause the redundant IBR 808 to change to the advantageous spectrum no longer being used by RF Link 852 with the failure condition.

FIG. 10 illustrates an exemplary embodiment of an IBR Antenna Array 648. FIG. 10 illustrates an antenna array having Q directive gain antennas 652 (i.e., where the number of antennas is greater than 1). In FIG. 10, the IBR Antenna Array 648 includes an IBR RF Switch Fabric 1012, RF interconnections 1004, a set of Front-ends 1008 and the directive gain antennas 652. The RF interconnections 1004 can be, for example, circuit board traces and/or coaxial cables. The RF interconnections 1004 connect the IBR RF Switch Fabric 1012 and the set of Front-ends 1008. Each Front-end 1008 is associated with an individual directive gain antenna 652, numbered consecutively from 1 to Q.

FIG. 10A illustrates an additional exemplary embodiment of an IBR Antenna Array 648, and comprises a block diagram of an IBR antenna array according to one embodiment of the invention relating to the use of dedicated transmission and reception antennas. In some IBR embodiments the embodiment of FIG. 10 may be replaced with the embodiments described in relation to FIG. 10A. For instance, such substitution may be made in use with either FDD, TDD, or even non-conventional duplexing systems. FIG. 10A illustrates an antenna array having $Q_R + Q_T$ directive gain antennas 652 (i.e., where the number of antennas is greater than 1). In FIG. 10A, the IBR Antenna Array 648 includes an IBR RF Switch Fabric 1012, RF interconnections 1004, a set of Front-ends 1009 and 1010 and the directive gain antennas 652. The RF interconnections 1004 can be, for example, circuit board traces and/or coaxial cables. The RF interconnections 1004 connect the IBR RF Switch Fabric 1012 and the set of Front-end Transmission Units 1009 and the set of Front-end Reception Units 1010. Each Front-end transmission unit 1009 is associated with an individual directive gain antenna 652, numbered con-

16

secutively from 1 to Q_T . Each Front-end reception unit 1010 is associated with an individual directive gain antenna 652, numbered consecutively from 1 to Q_R . The present embodiment may be used, for example, with the antenna array embodiments of FIG. 39, FIG. 41, FIG. 52, FIG. 53, or embodiments described elsewhere. Such dedicated transmission antennas are coupled to front-end transmission units 1009 and comprise antenna element 652.

In alternative embodiment, the IBR RF Switch fabric 1012 may be bypassed for the transmission signals when the number of dedicated transmission antennas and associated front-end transmission units (Q_T) is equal to the number of RF transmission signals RF-Tx-M (e.g. $Q_T = M$), resulting in directly coupling the IBR RF 632 transmissions to respective transmission front-end transmission units 1009. The dedicated reception antennas, comprising an antenna element 652 in some embodiments, are coupled to front-end reception units 1010, which in the present embodiment are coupled to the IBR RF Switch Fabric. In an additional alternative embodiment, the IBR RF Switch fabric 1012 may be bypassed for the reception signals when the number of dedicated reception antennas and associated front-end reception units (Q_R) is equal to the number of RF reception signals RF-Rx-N (e.g. $Q_R = N$), resulting in directly coupling the IBR RF 632 reception ports to respective front-end reception units 1010.

FIG. 11 illustrates an exemplary embodiment of the Front-end circuit 1008 of the IBR Antenna Array 648 of FIG. 10 for the case of TDD operation, and FIG. 12 illustrates an exemplary embodiment of the Front-end circuit 1008 of the IBR Antenna Array 648 of FIG. 10 for the case of FDD operation. The Front-end circuit 1008 of FIG. 11 includes a transmit power amplifier PA 1104, a receive low noise amplifier LNA 1108, SPDT switch 1112 and band-select filter 1116. The Front-end circuit 1008 of FIG. 12 includes a transmit power amplifier PA 1204, receive low noise amplifier LNA 1208, and duplexer filter 1212. These components of the Front-end circuit are substantially conventional components available in different form factors and performance capabilities from multiple commercial vendors.

As shown in FIGS. 11 and 12, each Front-end 1008 also includes an "Enable" input 1120, 1220 that causes substantially all active circuitry to power-down. Power-down techniques are well known. Power-down is advantageous for IBRs in which not all of the antennas are utilized at all times. It will be appreciated that alternative embodiments of the IBR Antenna Array may not utilize the "Enable" input 1120, 1220 or power-down feature. Furthermore, for embodiments with antenna arrays where some antenna elements are used only for transmit or only for receive, then certain Front-ends (not shown) may include only the transmit or only the receive paths of FIGS. 11 and 12 as appropriate.

FIG. 12A is a block diagram of a front-end transmission unit according to one embodiment of the invention relating to the use of dedicated transmission and reception antennas, and FIG. 12B is a block diagram of a front-end reception unit according to one embodiment of the invention relating to the use of dedicated transmission and reception antennas. As shown in FIGS. 12A and 12B, each Front-end 1008 also includes an "Enable" input 1225, 1230 that causes substantially all active circuitry to power-down, and any known power-down technique may be used. Power-down is advantageous for IBRs in which not all of the antennas are utilized at all times. It will be appreciated that alternative embodiments of the IBR Antenna Array may not utilize the "Enable" input 1225, 1230 or power-down feature. Furthermore, for some embodiments associated with FIG. 10A for example

17

(with antenna arrays where some antenna elements are used only for transmit or only for receive) then certain Front-ends may include only the transmit **1109** or only the receive paths **1010** of FIGS. **12A** and **12B** as appropriate. With respect to FIG. **12A**, Bandpass filter **1240** receives transmission signal RF-SW-Tx-qt, provides filtering and couples the signal to power amplifier **1104**, then to low pass filter **1050**. The output of the lowpass filter is then coupled to dedicated transmission antenna, which is comprised of directive antenna element **652**. With respect to FIG. **12B**, directive antenna element **652** is a dedicated receive only antenna and coupled to receive filter **1270**, when is in turn coupled to LNA **1208**. The resulting amplified receive signal is coupled to band pass filter **1260**, which provides output RF-SW-Rx-qr.

As described above, each Front-end (FE-q) corresponds to a particular directive gain antenna **652**. Each antenna **652** has a directivity gain Gq. For IBRs intended for fixed location street-level deployment with obstructed LOS between IBRs, whether in PTP or PMP configurations, each directive gain antenna **652** may use only moderate directivity compared to antennas in conventional PTP systems at a comparable RF transmission frequency. Based on measurements of path loss taken at street level at 2480 MHz in various locations in and nearby San Jose, Calif. during August and September 2010, IBR antennas should have a Gq of at least 6 dBi, and, in typical embodiments for operation between 2 GHz and 6 GHz RF transmission frequency, a Gq in the range of 10-18 dBi, wherein the radiation pattern in the elevation direction is typically less than 90° and nominally parallel to the local surface grade. At higher RF transmission frequencies, higher gain ranges of Gq are expected to be preferable. For example, Gq may be preferably 16-24 dBi for 20-40 GHz operation or 20-28 dBi for 60-90 GHz operation. In one particular embodiment, the directive gain antennas **652** are "patch" type antennas with Gq of about 13 dBi and nominally equal elevation and azimuth radiation patterns (about 40° each). Patch type antennas are advantageous because they can be realized using substantially conventional printed circuit board (PCB) fabrication techniques, which lowers cost and facilitates integration with Front-end circuit components and/or some or substantially all of the IBR RF Switch Fabric. However, many other antenna types, such as helical, horn, and slot, as well as microstrip antennas other than patch (such as reflected dipoles), and the like, may be used with the IBR Antenna Array. In an alternative embodiment, the directive gain antennas **652** are reflected dipoles with Gq of about 15 dBi (about 50° azimuth and 20° elevation). In many embodiments, the antenna elements are chosen with elevation angular patterns considerably less than azimuthal angular patterns.

In the IBR Antenna Array **648** illustrated in FIGS. **6**, **7** and **10**, the total number of individual antenna elements **652**, Q, is greater than or equal to the larger of the number of RF transmit chains **636**, M, and the number of RF receive chains **640**, N. In some embodiments, some or all of the antennas **652** may be split into pairs of polarization diverse antenna elements realized by either two separate feeds to a nominally single radiating element or by a pair of separate orthogonally oriented radiating elements. Such cross polarization antenna pairs enable either increased channel efficiency or enhanced signal diversity as described for the conventional PTP radio. The cross-polarization antenna pairs as well as any non-polarized antennas are also spatially diverse with respect to each other.

In some embodiments, certain antenna elements **652** may be configured with different antenna gain Gq and/or radiation patterns compared to others in the same IBR to provide pattern diversity.

18

In some embodiments, some antenna elements **652** may be oriented in different ways relative to others to achieve directional diversity. For example, FIG. **13** illustrates an IBR suitable for obstructed LOS PTP operation (or sector-limited PMP operation) in which spatial diversity (and optionally polarization diversity and/or pattern diversity) is utilized to the exclusion of directional diversity. As shown in FIG. **13**, all of the antenna elements **1304** are positioned on a front facet **1308** of the IBR. In FIG. **13**, the IBR **1300** includes eight antenna elements **1304** (Q=8). It will be appreciated that the IBR **1300** may include less than or more than eight antenna elements **1304**.

FIG. **14** illustrates another embodiment of an IBR **1400** where directional diversity is present. IBR **1400** includes the same number of antenna elements as the IBR **1300** shown in FIG. **13** (Q=8, or 16 if using cross-polarization feeds to all antenna elements). In FIG. **14**, the antenna elements **1404** are arranged on a front facet **1408** and two side facets **1412**. In FIG. **14**, the side facets **1412** are at a 45° angle in the azimuth relative to the front facet **1408**. It will be appreciated that this 45° angle is arbitrary and different angles are possible depending on the specific radiation patterns of the various antenna elements. Furthermore, the angle may be adjustable so that the side facets **1412** can vary in azimuth angle relative to the front facet between 0° to 90° (any value or range of values between 0° to 90°). Conventional electromechanical fabrication elements could also be used to make this side facing angle dynamically adjustable by, for example, the RRC **660** of FIG. **6** or the same in combination with the IBMS Agent **700** of FIG. **7**. Additionally, variations of the embodiment of FIG. **14** can use more than three facets at different angular spacing all within a nominal azimuthal range of approximately 180°, and the number of antenna elements **1404** may be less than or greater than Q=8. For example, in one embodiment, the antenna array includes four facets uniformly distributed in an azimuthal angular range across 160°.

FIG. **15** illustrates an IBR **1500** having an "omni-directional" (in the azimuth) array of antenna elements **1504**. In FIG. **15**, Q=16 antenna elements **1504** are uniformly distributed across all 360° of azimuth angles, amongst the eight facets **1508-1536**. Such an embodiment can be advantageous for propagation environments with severe obstructions between IBRs in a radio link or for an omni-directional common node at a point of aggregation (i.e. fiber POP) within a PMP deployment of IBRs. It will be appreciated that the IBR may have less than or more than eight facets, and that the number of antenna elements **1504** may be less than or greater than Q=16. It will also be appreciated that the antenna elements **1504** may be distributed non-uniformly across the facets.

With reference back to FIGS. **10-12**, the IBR RF Switch Fabric **1012** provides selectable RF connections between certain RF-Tx-m and/or RF-SW-Tx-q combinations and certain RF-Rx-n and/or RF-SW-Rx-q combinations. In an embodiment where Q=M=N, the IBR RF Switch Fabric **1012** can be parallel through connections or single-pole, single-throw (SPST) switches. In a maximally flexible embodiment where Q>Max (M, N) and any RF-Tx-m or RF-Rx-n can connect to any respective RF-SW-Tx-q or RF-SW-Rx-q, then a relatively complex cascade of individual switch blocks and a more extensive decoder logic may be required. Because each RF-Tx-m or RF-Rx-n can be readily interchangeable amongst their respective sets of signals at the digital baseband level, it is generally only necessary to connect any given RF-Tx-m or RF-Rx-n to a subset of the Front-ends **1008** roughly by the ratio respectively of Q/M or Q/N on average.

For example, if the IBR has $Q=8$ antenna elements and $M=N=4$, then $Q/M=Q/N=2$. Thus, any of the RF-Tx-m ($m=1, 2, 3, 4$) signals may be connectable to a pair of RF-SW-Tx-q signals, via a selectable RF connection including a SPDT switch (and similarly for RF-Rx-n to RF-SW-Rx-q). In this example, either RF-Tx-m and/or RF-Rx-n could connect via such a selectable RF connection to either one of the front-facing antenna elements or one of the side-facing antenna elements such that each RF signal has directional as well as spatial diversity options while allowing any two adjacent elements in the azimuth direction to both be selected. Similarly, for the IBR shown in FIG. 15, if the IBR has $Q=16$ (non-polarized) antenna elements and $M=N=4$, then any RF-Tx-n or RF-Rx-n signal could be oriented in one of four directions at 90° increments via a selectable RF connection including a single-pole, quadrature throw (SP4T) switch.

An alternative embodiment of the IBR RF Switch Fabric 1012 can also optionally connect, via a signal splitter, a particular RF signal (typically one of the RF-Tx-m signals) to multiple Front-ends 1008 and antenna elements 652 simultaneously. This may be advantageous in some IBR operational modes to provide an effectively broader transmit radiation pattern either in normal operation or specifically for certain channel estimation or broadcast signaling purposes. In context of the SPDT switch implementation in the example above for the IBR of FIG. 14, this would entail, if used for RF-Tx-m, the addition of another SPDT switch and three passive splitter/combiners as well as decoder logic for each antenna element pair.

In all of the foregoing descriptions of the IBR RF Switch Fabric 1012, substantially conventional components and RF board design structures as are well known can be used to physically implement such selectable RF connections. Alternatively, these selectable RF connections can also be realized by custom integrated circuits on commercially-available semiconductor technologies.

With reference back to FIGS. 6 and 7, the IBR RF 632 also includes transmit and receive chains 636, 640. Exemplary transmit and receive chains 636, 640 are shown in FIGS. 16 and 17 respectively. In one embodiment, as shown in FIG. 16, the transmit chain 636 takes a transmit chain input signal such as digital baseband quadrature signals I_{Tm} and Q_{Tm} , and then converts them to a transmit RF signal RF-Tx-m. Typically, each transmit chain Tx-m 636 includes at least two signal DACs, channel filters, a vector modulator, a programmable gain amplifier, an optional upconverter, and at least one synthesized LO. Similarly, as shown in FIG. 17, the receive chain 640 converts a receive RF signal RF-Rx-n to a receive chain output signal such as digital baseband quadrature signals I_{Rn} and Q_{Rn} . Typically, each receive chain Rx-n 640 includes an optional downconverter, a vector demodulator, an AGC amplifier, channel filters, at least two signal ADCs and at least one synthesized LO. A common synthesized LO can often be shared between pairs of Tx-m and Rx-n chains, or even amongst a plurality of such pairs, for TDD operation IBRs. Examples of commercially available components to implement the IBR RF chains include the AD935x family from Analog Devices, Inc. Numerous other substantially conventional components and RF board design structures are well known as alternatives for realizing the Tx-m and/or Rx-n chains whether for TDD or FDD operation of the IBR.

With reference back to FIGS. 6 and 7, the specific details of the IBR Modem 624 and IBR Channel MUX 628 depend somewhat on the specific modulation format(s) deployed by the IBR. In general, the IBR requires a modulation format suitable for a broadband channel subject to frequency-selective fading and multipath self-interference due to the desired

PHY data rates and ranges in obstructed LOS propagation environments. Many known modulation formats for such broadband channels are possible for the IBR. Two such modulation formats for the IBR are (1) Orthogonal Frequency Division Multiplexing (OFDM) and (2) Single-Carrier Frequency Domain Equalization (SC-FDE). Both modulation formats are well known, share common implementation elements, and have various advantages and disadvantages relative to each other.

As is well known, OFDM essentially converts the frequency-selective fading broadband channel into a parallel collection of flat-fading subchannels wherein the frequency spacing between subchannels is chosen to maintain orthogonality of their corresponding time domain waveforms. In OFDM, a block of information symbols is transmitted in parallel on a block of discrete frequency subchannels, each conveying one information symbol which can be efficiently channel multiplexed into the time domain by using an Inverse Discrete Fourier Transform (IDFT). A cyclic prefix of length in time greater than the dominant time delays associated with multi-path self-interference is then pre-pended to the IDFT output block by serially transmitting first in time a fraction of the IDFT output block time domain samples that are transmitted last. This length in time is also sometimes called a guard interval. The use of a cyclic prefix effectively converts a linear convolution of the transmitted block of symbols to a circular convolution such that the effects of inter-symbol interference (ISI) associated with multipath time delays can be largely eliminated at the OFDM receiver. At the OFDM receiver, the cyclic prefix is discarded and each time domain input block of symbols is demultiplexed back into frequency domain subchannels each conveying one information symbol by using a Discrete Fourier Transform (DFT). The transmission of a known training block of symbols within each OFDM PPDU enables the OFDM receiver to correct for carrier frequency offsets and determine a complex weighting coefficient for each frequency subchannel that can equalize the effects of frequency-selective relative gain and phase distortion within the propagation channel. Furthermore, transmission of known "pilot" sequences of symbols at certain predetermined subchannels within the transmit block enables the OFDM receiver to track such channel distortions and frequency offsets during reception of information symbol blocks the PPDU as well as provide a coherent reference for demodulation. Note that for those subchannels subjected to severe flat fading, as will occur inevitably in a broadband obstructed LOS propagation channel, the information within such subchannels cannot be directly demodulated. Thus to avoid a significant irreducible bit-error rate (BER) that would be unacceptable for most IBR applications, it is essential that either forward error correction (FEC) encoding be applied with a constraint length comparable to the number of bits per OFDM block of information symbols or with a combination of constraint length and interleaving depth such that related FEC encoded bits or symbols span substantially all of the OFDM block of information symbols.

In an SC-FDE transmitter, every block of information symbols, each mapped to the same single carrier frequency at a relatively high symbol rate, has a cyclic prefix prepended to it prior to transmission. Similar to OFDM, the cyclic prefix consists of a finite fraction of the modulated symbols with a length in time greater than the dominant time delays associated with multipath self-interference wherein such modulated symbols are identical to those to be transmitted last in time for each block. Analogously to OFDM, this cyclic prefix effectively converts a linear convolution of the transmitted block of symbols to a circular convolution such that inter-block inter-

ference (IBI) due to multipath can be largely eliminated at the SC-FDE receiver. The SC-FDE receiver is similar to the OFDM receiver in that a cyclic prefix remover discards a cyclic prefix for each block of information symbols and the remaining sampled signals are decomposed into a set of frequency subchannels collectively representing the IBI-free block of symbols using a DFT. Based on a set of complex weighting coefficients, one for each frequency sub-channel, as usually determined from a known training block of symbols within each SC-FDE PPDU, the broadband channel induced relative distortions of amplitude and phase for each frequency sub-channel are then corrected in the Frequency Domain Equalizer (FDE). In contrast to OFDM where FDE-corrected subchannel estimates can be directly demapped as individual information symbols, in SC-FDE the FDE-corrected frequency domain subchannel estimates are then remultiplexed into a channel equalized single-carrier stream of information symbol estimates using an IDFT so that such information symbol estimates can be subsequently demodulated.

Embodiments of the IBR may use Quadrature Amplitude Modulation (QAM) to map groups of information data bits to a symbol including an I (or “real”) component and a Q (or “imaginary”) component. These symbols (i.e., symbols that include I and Q components) are typically referred to as “complex” symbols. Such “complex” symbols may be multiplexed or demultiplexed by an IDFT or DFT respectively implemented by structures executing a complex Inverse Fast Fourier Transform (IFFT) or complex Fast Fourier Transform (FFT). In IBR embodiments, references to IDFT or DFT herein assume that such transforms will typically be implemented by structures executing an IFFT or FFT respectively. Note also that the cyclic prefix described above can also be implemented as a cyclic postfix for either OFDM or SC-FDE with equivalent performance. For SC-FDE, some re-ordering of samples at the receiver after removal during the guard interval may be required if a cyclic postfix is used. It will be appreciated however that techniques other than QAM for modulation mapping may also be used.

With reference again to FIGS. 6 and 7, the specific details of the IBR Modem 624 and IBR Channel MUX 628 also depend somewhat on the specific antenna array signal processing format(s) deployed by the IBR. In general, the IBR utilizes multiple antennas and transmit and/or receive chains which can be utilized advantageously by several well-known baseband signal processing techniques that exploit multipath broadband channel propagation. Such techniques include Multiple-Input, Multiple-Output (MIMO), MIMO Spatial Multiplexing (MIMO-SM), beamforming (BF), maximal ratio combining (MRC), and Space Division Multiple Access (SDMA).

In general, any configuration where a transmitter that has multiple transmit antennas (or “inputs” to the propagation channel) communicates with a receiver that has multiple receive antennas (or “outputs” from the propagation channel) can technically be described as MIMO. However, typically the term MIMO is used in the context of spatially multiplexing multiple encoded and modulated information streams from multiple transmit antennas into a multipath channel for receipt at multiple receive antennas wherein inversion of channel transfer matrix, usually determined by a known training block associated with each PPDU, enables separation and demodulation/decoding of each information stream—a process called MIMO-SM. Various embodiments of the IBR, as described herein, advantageously use other types of antenna diversity such as directional or polarization diversity to real-

ize MIMO-SM performance even in propagation environments where spatial separation alone may be inadequate for conventional MIMO-SM.

For a given encoded and modulated information stream, BF or MRC can be utilized at either the transmitter or the receiver (or both) to improve either the signal to interference and noise ratio (SINR) or the signal to noise ratio (SNR). For example, BF or MRC optimally combine the information signal received from the multiple antennas or split the information stream signal as transmitted by the multiple antennas. Numerous algorithms for determining the optimal weighting criteria amongst the multiple antennas, usually as a function of frequency within a frequency-selective multipath broadband channel, are well known.

SDMA allows an Aggregation End IBR (AE-IBR) in a PMP configuration (see FIG. 5) to transmit to or receive from multiple Remote End IBRs (RE-IBRs) simultaneously using parallel encoded and modulated information streams each associated with multiple antenna and transmit/receive chain combinations wherein stream-specific BF maximizes signal quality for a given stream and IBR pair, while minimizing interference to other streams and IBR pairs. To the extent that multiple RE-IBRs are separable in space, or another antenna array characteristic such as polarization or direction, then significant increases in overall spectrum efficiency of a PMP system are possible using SDMA. As with BF (and MRC), numerous algorithms for computing the optimal weighting criteria, such as Eigen Beam Forming (EBF), amongst multiple antennas are well known.

In view of the foregoing exemplary modulation format and antenna array processing format alternatives for the IBR, exemplary embodiments of the IBR Modem 624 and IBR Channel MUX 628 are described with reference to FIGS. 18-30.

FIG. 18 shows an exemplary embodiment of the IBR Modem 624. The IBR modem 624 includes a Tx PLCP generator 1804, an Rx PLCP analyzer 1808, multiple modulator cores 1812 and multiple demodulator cores 1816. FIG. 18 is divided functionally into a transmit path and a receive path. The transmit path includes the TxPLCP 1804 and the modulator cores 1812, and the receive path includes the Rx PLCP analyzer 1808 and the demodulator cores 1816.

The transmit path of the IBR Modem 624 includes a total of J_{mod} “Modulator Cores” 1812, each denoted as Modulator Core j wherein $j=1, 2, \dots, J_{\text{mod}}$. The Tx PLCP generator 1804 provides transmit data interface streams, $Tx-j$, to each Modulator Core j 1812 such that A_j total vector outputs of mapped transmit symbol streams (denoted by the “+” on such l, Q connections) are generated from each Modulator Core j 1812. This results in a total number of transmit symbol streams, $K=A_1+A_2+\dots+A_{J_{\text{mod}}}$, each in vector format $(I_{TS}, Q_{TS})_k$ from $k=1$ to K that are connected to the transmit path of the IBR Channel MUX 628.

Similarly, the receive path of the IBR Modem 624 includes a total of J_{dem} “Demodulator Cores” 1816 each denoted as Demodulator Core j 1816 wherein $j=1, 2, \dots, J_{\text{dem}}$. The IBR Channel MUX receive path provides $L=B_1+B_2+\dots+B_{J_{\text{dem}}}$ vector format receive symbol streams $(I_{RS}, Q_{RS})_l$ for $l=1$ to L that are input as B_j vector streams per each Demodulator Core j 1816 to produce the receive data interface stream $Rx-j$.

In PTP IBR configurations where $J_{\text{mod}}=J_{\text{dem}}$, usually $A_j=B_j$. However, for a PTP IBR where probing capability is present, $A_j=B_j$ only for $j=1$ to J_{mod} in cases where $J_{\text{dem}}>J_{\text{mod}}$. In PMP IBR configurations, it may happen that $A_j \neq B_j$; even if $J_{\text{mod}}=J_{\text{dem}}$.

An exemplary embodiment of a Modulator Core j 1812 is illustrated in FIG. 19. As shown in FIG. 19, this exemplary

23

modulator core **1812** includes a scrambler **1908**, encoder **1912**, stream parser **1916**, multiple optional interleavers **1920**, multiple symbol groupers **1924** and multiple symbol mappers **1928**.

Typically, the data from the Tx PLCP generator **1804**, Tx-j **1904**, is scrambled at the scrambler **1908** and then passed to the encoder **1912**. The encoder **1912** provides FEC encoding and in some types of encoders also effectively interleaves the encoded data. Exemplary FEC encoder types for IBRs include convolutional, turbo and low density parity check (LDPC). The encoded data is passed to the Stream Parser **1916**. The Stream Parser **1916** demultiplexes the encoded data into Aj streams. Each encoded data stream is then interleaved if necessary at the optional interleavers **1920** such that the greater of the FEC encoder constraint length and the interleaving depth approximates the total number of bits per transmitted block of symbols in either OFDM or SC-FDE. Example interleaver types include block (or “row/column”) and convolutional. Such interleaved and/or encoded data in each stream is then grouped at symbol groupers **1924** based on the number of encoded bits per symbol. The groups of bits pass to the symbol mapper **1928**, which converts each group of encoded bits into a vector representation (I, Q) within a constellation and then provides an output as a transmit symbol stream **1932**. An exemplary technique for mapping encoded bits to a vector symbol is QAM optionally with Gray coding.

FIG. **20** illustrates an exemplary embodiment of a Demodulator Core j **1816** compatible with the exemplary Modulatory Core j **1812** of FIG. **19**. As shown in FIG. **20**, the demodulator core **1816** includes a descrambler **2008**, decoder **2012**, stream MUX **2016**, multiple optional deinterleavers **2020**, and multiple soft decision symbol demappers **2024**.

The Demodulator Core j **1816** at the highest level can be described as performing essentially the reverse operations of those performed in the Modulator Core j **1812** of FIG. **19**. A key difference though is that the vector representation (I, Q) receive symbol streams **2032** input to each Demodulator Core j **1816** are only estimates which may be corrupted due to channel or receiver impairments such as multipath self-interference, Gaussian noise, co-channel interference, carrier frequency offset (CFO), distortion through channel filters at Tx and/or Rx, non-linearity in the Tx and/or Rx chains and Front-ends, or phase noise in either Tx and/or Rx local oscillators. Thus, the Demodulator Core j **1816** may use the soft decision symbol demapper **2024**, which estimates the likelihood that a received symbol or bit has a particular value using, for example, a known technique such as Log-Likelihood Ratio (LLR). For example, if each data bit out of the demapper **2024** had a soft-decision representation (or “metric”) of 8 bits, then a value of 0 would represent a data bit of 0, a value of 1-20 would indicate most likely a data bit of 0, a value of 21-125 would indicate more likely a data bit of 0 than a data bit of 1, a value of 126-129 would indicate near uncertainty of the data bit as either 0 or 1, a value of 130-234 would indicate more likely a data bit of 1, a value of 235-254 would indicate most likely a data bit of 1, and a value of 255 would represent a data bit of 1. These soft-decision metrics can then be deinterleaved if applicable at optional deinterleavers **2020** and stream multiplexed at the Stream MUX **2016**, and then supplied to the decoder (deinterleaver) **2012** as a sequence of soft-decision metrics estimating the originally encoded (and possibly interleaved) bit stream. Decoder types are matched to encoder types as is well known. Techniques such as iterative decoding within the IBR modem or combination of IBR modem and IBR Channel MUX are also known and can be used in some embodiments.

24

FIGS. **19** and **20** are examples of implementations of the modulator and demodulator cores **1812**, **1816**, respectively, in the IBR Modem **624**. It will be appreciated that the order of the Scrambler **1908** and Encoder **1912** can be reversed (and hence the Decoder **2012** and Descrambler **2008** could be reversed). Also, the Stream Parser **1916** can divide a transmit data interface stream Tx-j sequentially into multiple encoder/scrambler combinations and hence the corresponding Stream MUX **2016** would combine multiple Decoder outputs. FIGS. **21** and **22** illustrate exemplary alternatives with the above recited differences and without showing optional separate interleavers **1920** and deinterleavers **2020**.

Although each element is illustrated separately in FIGS. **19-22**, the elements are not required to be distinct physical structures. It is feasible to time multiplex either the entire Modulator or Demodulator Cores **1816**, **1820** or constituent elements such as the Encoder **1912** or Decoder **2012** across multiple streams. Furthermore, in TDD operation, it can be advantageous to time multiplex constituent elements such as the Encoder **1912** and/or Decoder **2012** even during times where the opposite transmit/receive path is over the air compared to the time multiplexed element. For example, by buffering receive symbol stream samples from the receive path of the IBR Channel MUX **628** during the receive portion of TDD operation, some of the streams can be decoded in a common Decoder **2012** even when the IBR has changed over to transmit operation.

Note further that the exemplary embodiments of FIGS. **19-22** for the Modulator and Demodulator Cores of the IBR Modem are all suitable for either OFDM or SC-FDE operation.

FIG. **23** (FIGS. **23A** and **23B** together) illustrates an exemplary embodiment of the IBR Channel MUX **628**. FIG. **23** is divided functionally into a transmit path (FIG. **23A**) and a receive path (FIG. **23B**). In FIG. **23A**, the transmit path includes multiple block assemblers **2304**, multiple transmit channel equalizers (CEs) **2308**, multiple transmit MUXs **2312**, multiple cyclic prefix addition (CPA) and block serializers **2316**, multiple transmit digital front ends (DFEs) **2320**, an optional preamble samples library **2324**, and a pilot (and optionally preamble) symbol library **2328**. In FIG. **23B**, the receive path includes an acquisition/synchronization calculator **2336**, multiple receive digital front ends (DFEs) **2340**, multiple cyclic prefix removal (CPR) and block assemblers **2344**, multiple complex DFTs **2348**, multiple receive channel equalizers **2352** and multiple receive MUXs **2356**. The IBR Channel MUX **628** also includes a channel equalizer coefficients generator **2332** (see FIG. **23A**). The channel equalizer coefficients generator **2332** may be used by both transmit and receive paths or only the receive path depending on the specific embodiment. In the receive path, the IBR Channel MUX **628** is frequency selective to enable operation in obstructed LOS propagation environments with frequency selective fading within the channel bandwidth as well as operation in predominantly unobstructed LOS.

In the transmit path, each transmit symbol stream of mapped symbols $(I_{TS}, Q_{TS})_k$ from $k=1$ to K **2360** is connected to a respective Tx Block Assembler k **2304** as shown in FIG. **23A**. For OFDM modulation, each Tx block Assembler k **2304** is equivalent to a serial to parallel buffer that also places “running” pilot symbols in pre-determined locations corresponding to pilot subchannels. In the exemplary embodiment of FIG. **23A**, such pilot symbols are supplied to each Tx Block Assembler k **2304** by a Pilot Symbol Library **2328**. Alternatively, pilot symbols are selectively injected into the data stream at predetermined points with a Modulator Core **1812** similar to FIG. **19** or **21** but with a “Pilot Data Library”

25

(not shown) providing grouped data to a selectable buffer or multiplexer (not shown) between the Symbol Grouper **1924** and Symbol Mapper **1928** of either FIG. **19** or **21**.

For SC-FDE modulation with no frequency selective channel equalization, the Tx Block Assembler **k 2304** would, as an exemplary embodiment, be a simple serial to parallel buffer that places pilot symbols in pre-determined symbol sequence positions either from the Pilot Symbol Library **2328** as shown, or alternatively from a "Pilot Data Library" in the Modulator Core as described above. If frequency selective channel equalization in Tx for SC-FDE is used, then in addition to the above, the Tx Block Assembler **k 2304** would also include a Complex DFT that follows the serial to parallel buffer. In this DFT structure version of SC-FDE (also known as DFT pre-coding SC-DFE), it is also possible to insert the pilot symbols as "pseudo-subchannels" after the DFT operation instead of as time domain symbols before the DFT.

In alternative embodiments that utilize frequency selective channel equalization in Tx for SC-FDE, the transmit channel equalizer **2308** may utilize a time domain approach as shown in FIG. **24A**. The embodiment shown in FIG. **24A** does not require the Complex DFT as part of the Tx Block Assembler **k 2304** shown in FIG. **23A** because the transmit equalization is performed in the time domain, obviating the need for the DFT or the inverse DFT (IDFT).

After block assembly, the blocks of mapped symbols (each a vector of (I, Q) constellation points) are typically supplied to each of M transmit channel equalizers **2308** ("Tx-CE-m" for m=1 to M in FIG. **23A**). Each Tx-CE-m **2308** applies K blocks of amplitude and/or phase weights (each amplitude phase weight may be represented in Cartesian form as a "complex" weight with a "real" and "imaginary" sub-component values) respectively to each of the K blocks of mapped vector symbols.

An exemplary embodiment of a transmit channel equalizer Tx-CE-m **2308** is shown in FIG. **24**. The blocks of mapped vector symbols are $\overrightarrow{TxBlock_k}$ wherein each symbol element within the vector $\overrightarrow{TxBlock_k}$ has an (I, Q) or, relative to FIG. **24**, "complex" value (I representing a "real" sub-component value, Q representing an "imaginary" sub-component value). Each $\overrightarrow{TxBlock_k}$ vector of complex symbols is then complex multiplied by a respective transmit weight vector $\overrightarrow{WT_{m,k}}$ as shown in FIG. **24**. The channel equalized vector of symbols

for each transmit chain m is then $\overrightarrow{TxBlockEq_m}$ (the transmit chain channel-equalized symbols) which represents the complex summation of the symbol by symbol output vectors of each respective weighting operation. Note that FIG. **24** depicts an exemplary implementation of transmit channel equalizer TX-CE-m **2308** based on vector complex multiplication and summation. Such an approach is easily implemented using either generic function calls or custom software modules executing on an embedded processor (or Digital Signal Processor—"DSP"). However, numerous other implementation techniques are also known such as dedicated logic circuits for complex multiplication and/or summation, a set of four scalar multiplier circuits and two scalar adders, or other logic circuits such as combinatorial gates (i.e. OR, NOR, AND, NAND, XOR, XNOR, etc.), multiplexers, shift registers, etc. that can produce an equivalent result.

An alternative exemplary embodiment of a transmit channel equalizer Tx-CE-m **2308** is shown in FIG. **24A**. The transmit channel equalizer Tx-CE-m **2308** allows for time domain based processing. The transmit channel equalizer Tx-CE-m **2308** is similar to the TX-CE-m shown in FIG. **24**,

26

but not need a DFT in the Tx Block Assembler **k 2304** or an IDFT within the Tx-Mux-m **2312**. As described in relation to FIG. **24**, the blocks of mapped vector symbols are

$\overrightarrow{TxBlock_k}$ wherein each symbol element within the vector $\overrightarrow{TxBlock_k}$ has an (I, Q) or, relative to FIG. **24**, "complex" value (I representing a "real" sub-component value, Q representing an "imaginary" sub-component value). In the embodiment of FIG. **24A**, the processing of these complex values is performed differently, having a distinct advantage in complexity for a lower number of channel equalization impulse response length relative to the traditional approach for frequency domain equalization as described in relation to

FIG. **24**. Each $\overrightarrow{TxBlock_k}$ vector of complex symbols passed through a complex finite impulse response filter (FIR) **2401**. Such structures and DSP processing implementations are well known in the art and comprise a mathematical convolution with the FIR filter weights. Such weights result in a desired impulse response and are provided to the FIR **2401** by the transmit weight vector $\overrightarrow{WT_{m,k}}$ as shown in FIG. **24A**. A circular summer **2405** is also used to convert the linear convolution of **2401** to a circular convolution. The circular convolution retains the original number of K symbols of the

$\overrightarrow{TxBlockEq_m}$, rather than expanding them by one less than the length of the FIR filter weights (as is the case with linear convolution). The channel equalized vector of symbols for

each transmit chain m is then $\overrightarrow{TxBlockEq_m}$ (the transmit chain channel-equalized symbols) which represents the complex summation of the symbol by symbol output vectors of each respective circular convolution operation between each $\overrightarrow{TxBlock_k}$ and the transmit weight vector $\overrightarrow{WT_{m,k}}$.

The circular FIR filtering is described in detail with reference to FIGS. **24B-24I**. FIG. **24B** is a plot of an exemplary transmit equalizer impulse response for use with the Tx-CE-m having a complex circular FIR capability. FIG. **24C** is a plot of an exemplary transmit equalizer frequency response for use with the Tx-CE-m having the complex circular FIR capability.

FIGS. **24D**, **E**, and **F** are presented together on the same drawing sheet for purposes of discussion. FIG. **24D** is equivalent to FIG. **24B**. FIG. **24E** is a plot of an exemplary 16

QAM-10 symbol Transmit Block ($\overrightarrow{TxBlock_k}$), where K=10 (K is the number symbols contained within an example transmit block). It will be appreciated that K is typically much larger, but a shorter block size of K=10 is utilized to easily illustrate the processing with the Tx-CE-m having the complex circular FIR capability. FIG. **24F** shows the result of a linear convolution of the exemplary 16 QAM-10 symbol

Transmit Block ($\overrightarrow{TxBlock_k}$) with the example complex impulse response. A linear convolution between K samples and WT samples results in K+WT-1 value. However, an expansion in the number of transmit symbols is undesirable and does not provide for the equivalent processing of a frequency domain implementation of FDE processing. Taking the WT-1 last samples of the resulting linear convolution output, and summing them respectively with the first WT-1 samples of the linear convolution output results in an equivalent circular convolution. Such processing may be performed in the circular summer **2405** shown in FIG. **24A**.

FIG. **24G** is a plot of the result of the FIR based circular convolution utilizing the linear FIR **2401** and the circular

27

summer (2405) of the exemplary 16 QAM-10 symbol Transmit Block ($\overline{TxBlock_k}$) with the example complex impulse response. FIG. 24H is a plot of the result of a frequency domain based circular convolution of the exemplary 16

QAM-10 symbol Transmit Block ($\overline{TxBlock_k}$) with the example complex impulse response. FIG. 24G and FIG. 24H show the resulting circular convolutional processing output as performed respectively by the circular FIR of FIG. 24A, and traditional frequency domain circular convolutional based processing as described in relation to the frequency selective transmit channel equalization, when applied to SC-FDE. FIG. 24I is a magnitude plot of the result of both the circular FIR and the frequency domain based circular convolutions of the exemplary 16 QAM-10 symbol Transmit Block ($\overline{TxBlock_k}$) with the example complex impulse response for comparison, showing that these values are identical. Such use of the time domain processing for frequency selective transmit equalization eliminates the requirement for both a DFT and an IDFT, resulting in reduced processing requirements for transmit weight vectors of moderate length.

Note further that for embodiments with either OFDM or SC-FDE modulation and no frequency selective channel equalization, $\overline{WT}_{m,k}$ would typically be composed of a block of identical transmit weights applied equally to all mapped symbols for a given stream k. In some OFDM or SC-FDE embodiments where $K=M$, no equalization or weighting amongst transmit streams and chains may be desired such that each Tx-CE-m 2308 in FIG. 23A (or FIG. 24) is equivalent to a through connection mapping each $\overline{TxBlock_k}$ directly to each $\overline{TxBlockEq_m}$ for each $k=m$.

For each transmit chain Tx-m, the channel equalized vector of symbols for each transmit chain $\overline{TxBlockEq_m}$ is supplied to Tx-Mux-m 2312, where $m=1$ to M , as shown in FIG. 23A. In the case of OFDM modulation or SC-FDE modulation with frequency-selective channel equalization utilizing the TX-EQ-m 2308 of FIG. 24, a DFT is required within the TX Block Assembler k 2304, and each TX-Mux-m 2312 includes an IDFT to transform the successive frequency domain symbols into a block of time domain samples. Note that in the case of embodiments of the TX-EQ-m of FIG. 23A (the time domain frequency selective equalization) no such IDFT is required. For SC-FDE modulation with no frequency selective channel equalization, and hence no DFT pre-coding within each Tx Block Assembler k 2304, then $\overline{TxBlockEq_m}$ is already effectively a block of transmit chain time domain samples and no such IDFT is required.

With reference back to FIG. 23A, each respective Tx-Mux-m 2312 is followed by a CPA and Block Serializer-m 2316. "CPA" refers to "Cyclic Prefix Addition"—cyclic prefix was described above—which is performed by the cyclic prefix adder within each CPA and Block Serializer-m 2316. Each CPA and Block Serializer-m 2316, via a cyclic prefix adder, typically replicates a set of samples from the end of the block and prepends these samples to the beginning of a cyclically-extended block or includes logic to read out these samples in a way that produces an equivalent result. The CPA and Block Serializer-m 2316, via a block serializer, then effectively performs a parallel to serial conversion of the extended block into a sequence of cyclically-extended (I, Q) transmit chain time domain samples.

28

The cyclically-extended block of (I, Q) time domain samples for each transmit chain Tx-m are then supplied to a respective Tx-DFE-m 2320 (DFE refers to Digital Front End). Each Tx-DFE-m 2320 performs a variety of possible time domain signal processing operations that may be specific to particular IBR embodiments. Common DFE processes in the transmit path include digital up-conversion of the sampling rate, DC offset calibration for each Tx-m transmit chain, digital pre-distortion to account for non-linearities in the analog transmit path, pulse-shaping filtering, crest factor or peak to average power ratio (PAPR) reduction, or frequency shifting to an "intermediate Frequency" (IF) suitable for analog up conversion to RF-Tx-m in the Tx-m transmit chain (as opposed to the baseband $(I_T, Q_T)_m$ interface illustrated in FIG. 16). Except for the last listed option of digital IF up conversion, the output (a transmit chain input signal) of each respective Tx-DFE-m 2320 is typically a sequence of cyclically-extended blocks of calibrated, compensated and filtered baseband symbol samples $(I_T, Q_T)_m$.

In the receive path of the exemplary IBR Channel MUX 628, downconverted and amplified samples from I and Q ADCs in each Rx-n receive chain (a receive chain output signal) are passed to respective Rx-DFE-n 2340. Although the receive path of the IBR generally follows the logic of reversing the operations of the transmit path, the details are considerably different in the IBR Channel MUX because in the receive path the samples are corrupted by channel propagation impairments and arriving at initially unknown times. In view of this the receive path Digital Front Ends of the IBR Channel MUX shown in FIG. 23B, Rx-DFE-n 2340, perform several time domain digital operations typically including matched filtering, rotating, sampling rate downconversion, and Rx-n chain calibration. At minimum, the Rx-DFE-n 2340 is used in the detection of timing synchronization information in the optional Preamble of the L PPDU-L transmissions from the transmitting IBR.

FIG. 26 illustrates this Preamble in the time domain. Note that the optional Preamble for the l-th Tx-path in a transmitting IBR could be either generated from a symbol library (as indicated optionally in Pilot (And Preamble) Symbol Library 2328 of FIG. 23) or read directly from a Preamble Samples Library 2324 (also indicated optionally in FIG. 23).

Referring again to FIG. 23B, the Acquisition/Synchronization detector 2336 acquires a timing reference, for example, from detection of one or more of the up to L optional preambles, one or more of the blocks of Training Pilots, an optional known pattern within a PLCP header, and/or pilot symbols within blocks of transmit symbol streams. The Acquisition/Synchronization detector 2336 provides a feedback signal to each Rx-DFE-n 2340 that enables, for example, symbol rotation and alignment in the time domain (typically in the RX-DFE-n using a CORDIC algorithm). The Acquisition/Synchronization detector 2336 further provides symbol boundary and block boundary timing references (not shown) for use by various elements of the IBR Channel MUX 628 and the IBR Modem 624.

Such corrected receive symbol samples are then supplied to a respective CPR and Block Assembler-n 2344. "CPR" means "cyclic prefix removal." The CPR and Block Assembler-n 2344 effectively discards the number of received samples at beginning of each block corresponding to the number of cyclic prefix samples prepended to each block in the transmit path of the transmitting IBR. The remaining samples are serial to parallel buffered into a single block (or an equivalent operation) suitable for decomposition into receive chain frequency domain subchannels in each respective Complex DFT-n 2348.

As shown in FIG. 23B, the sets of receive chain frequency domain subchannel samples output from the N Complex DFT-n **2348** are collectively supplied to L frequency domain receive channel equalizers **2352**, each a Rx-CE-l. FIG. 25 illustrates an exemplary embodiment of Rx-CE-l **2352**. The receive channel equalizer Rx-CE-l **2352** is analogous to the transmit channel equalizer Tx-CE-m **2308** shown in FIG. 24. Although implementation details of each Rx-CE-l **2352** can vary considerably just as described above for TX-CE-m **2308**, essentially the frequency domain receive vectors of I and Q

samples for each block, $\overrightarrow{RxBlock_n}$, are respectively complex multiplied by complex receive weights $\overrightarrow{WR_{l,n}}$ and then complex summed to produce a set of channel-equalized frequency domain estimates, $\overrightarrow{RxBlockEq_l}$, representative of the original transmitted block corresponding to stream l.

The task of producing such complex receive weight vectors $\overrightarrow{WR_{l,n}}$ that allow each Rx stream l to be separated from the myriad of signals, both desired and undesired, received via N ($N \geq L$) receive chains is performed by the Channel Equalizer Coefficients Generator **2332** of FIG. 23A. There are many known algorithms for generating the appropriate frequency selective complex receive weight vectors for the IBR Channel MUX as shown in FIG. 23 wherein a known "Training Block" of pilot symbols for each stream l is used to determine the receive weight vectors $\overrightarrow{WR_{l,n}}$ that allow linear detection of

each $\overrightarrow{RxBlockEq_l}$ vector in view of the actual receive chain frequency domain subchannel samples observed for the N receive chains during a reception period corresponding to transmission of the Training Block (or block of Training Pilots). In general, a Channel Equalizer Coefficients Generator **2332** that is based on these algorithms calculates the various weights by comparing the actual receive chain frequency domain subchannel samples observed during the reception of the Training Block with the expected frequency domain subchannel samples corresponding to a particular known Training Block for any given transmit stream. Such algorithms include Zero Forcing (ZF), Maximal Ratio Combining (MRC) and Minimum Mean Square Error (MMSE). Alternatively, other known algorithms for non-linear detection of each $\overrightarrow{RxBlockEq_l}$ are also known, such as Successive Interference Cancellation (SIC) in combination with ZF, MRC or MMSE, V-BLAST (an acronym for "vertical Bell Labs Layered Space-Time) with ZF or MMSE, or Maximal Likelihood (ML).

The Channel Equalizer Coefficients Generator **2332** also supplies the complex transmit weight vectors $\overrightarrow{WT_{m,k}}$ used for transmit channel equalization within the exemplary IBR Channel MUX **628**. In an ideal PTP TDD configuration where $K=L$, $M=N$, and no other co-channel interference beyond the multiple transmit streams, for frequency domain transmit channel equalization, a straightforward alternative is to derive $\overrightarrow{WT_{m,k}}$ directly from the computed $\overrightarrow{WR_{l,n}}$ of the previous superframe. However, this can be sub-optimal for situations in which a co-channel interferer, such as for example either another IBR PTP or PMP link in the same vicinity or a conventional PTP link nearby, affects the received signals at the N receive chains of the AE-IBR differently from those at the RE-IBR. An alternative is to calculate $\overrightarrow{WR_{l,n}}$ using both MMSE (which will maximize SINR at the receiver) and MRC (which will maximize SNR, or effec-

tively maximize signal power, at the receiver) and then derive $\overrightarrow{WT_{m,k}}$ from the computed, but unused in receive channel equalization, $\overrightarrow{WR_{l,n}}$ for MRC. Note that for SC-FDE with no frequency selective channel equalization, then the constant $\overrightarrow{WT_{m,k}}$ values applied to all symbols in a $\overrightarrow{TxBlock_k}$ can be

blended using known algorithms from vector $\overrightarrow{WT_{m,k}}$ values derived for MRC. This allows such an SC-FDE PTP transmitter to improve the overall signal quality at the other receiver while allowing the other receiver to equalize both interference and frequency selective fading effects. In view of the foregoing, it is advantageous for the PMP AE-IBR to use either OFDM or SC-FDE with transmit frequency selective channel equalization such that the Channel Equalizer Coefficients Generator can compute MIMO-SDMA complex weights, using for example known EBF or Space Time Adaptive Processing (STAP) algorithms, to minimize multi-stream interference to RE-IBRs that beneficially receive only a subset of the transmitted streams. If SC-FDE without frequency selective transmit equalization is used, there is still a benefit in

using scalar $\overrightarrow{WT_{m,k}}$ transmit weights at the AE-IBR derived for streams to a given RE-IBR from previous superframe MRC at the AE-IBR receiver but some additional signal separation such as time and/or frequency may be required to direct data to specific RE-IBRs.

With reference back to FIG. 23B, each respective channel equalized and interference cancelled symbol samples stream

vector $\overrightarrow{RxBlockEq_l}$ (or set of channel-equalized frequency domain estimates) is then supplied to an Rx-MUX-l **2356**. For OFDM, this is typically a parallel to serial converter that removes the pilot subchannel symbol samples which are not needed in the IBR Modem. For SC-FDE, each Rx-MUX-l **2356** includes a Complex IDFT and a parallel to serial converter that removes the time domain symbol samples at the pilot symbol sequence positions. Note though that for IBR embodiments where the transmitter uses SC-FDE with frequency selective equalization and inserts the pilot symbols as "pseudo-subchannels" (see above), each Rx-MUX-l also needs to discard the pilot subchannels prior to the Complex IDFT transformation back to a time domain symbol samples stream. The output of each Rx-MUX-l is a receive symbol stream l that is provided to a respective input to a demodulator core j.

FIG. 26 illustrates block by block generation of a given transmit symbol stream l in the transmitter IBR as combined into a transmit chain m. The designator "l" is used in FIG. 26 instead of "k" for the transmitter to denote the generation of what becomes the "l-th" receive symbol stream being demodulated at the receiver. Per the above description of FIG. 23, optionally every PPDU-l can start with a Preamble of known samples that enables rapid signal acquisition and coarse synchronization. However, for some embodiments of the IBR in either FDD or in TDD with fixed and short super-frame timing, it is not necessary to perform such acquisition and synchronization from the start of every PPDU. Thus, overhead can be advantageously minimized by either discarding the Preamble from the PPDU-l generation process, at least for subsequent PDUs after startup, or effectively combining the desired characteristics of the Preamble into the Training Block via the Pilots-l. This approach can also be used with a Channel Sense Multiple Access (CSMA) approach to MAC scheduling to the extent that a timing master in the system, usually the AE-IBR, maintains sufficiently frequent timing synchronization transmissions.

31

In some embodiments, the PLCP Header for stream-1 is matched to a particular block such that the modulation and coding scheme (MCS) of such PLCP Header block (see, for example, Block 1 of FIG. 26) is always predetermined and usually chosen to deliver the PLCP header with higher reliability than the data payload blocks to follow. In typical embodiments, the MCS includes an index number that conveys information regarding parameters used in the modulation mapping and the forward error correction encoding processes. In such embodiments, the PLCP Header typically conveys at least the information needed to inform the receiver at the receiving-end IBR of the MCS used for the data payload blocks to follow. In other embodiments, the MCS or information indicating an increment, decrement or no change amongst a list of possible MCS values is conveyed within Block 0 of FIG. 26 by, for example, adjusting the modulation of certain pilot symbols (or pilot subchannels). In the case where Block 1 of FIG. 26 can carry more data than just the PLCP Header alone, then the IBRs may allocate PPDU payload to Block 1 even though the bits may be transferred at a different MCS than that of subsequent Blocks.

In some embodiments, each Block 1 through f (or Block 1 through r) may correspond to the output of a block encoder and/or block interleaver of FIG. 19 or 21. “f” denotes a forward link transmitted from an AE-IBR to an RE-IBR and “r” denotes a reverse link transmitted from an RE-IBR to an AE-IBR. Furthermore, each Block may include error detection bits (such as an FCS) or error correction bits (such as for a Reed-Solomon or a Reed-Muller outer encoder) that can be used either for an Automatic Repeat reQuest (ARQ) re-transmission protocol of certain blocks or for input to an RLC and/or RRC entity at either the AE-IBR or RE-IBR to optimize future link parameters or resource selections.

FIG. 27 illustrates an exemplary embodiment of the Tx PLCP generator 1804 of the IBR modem of FIG. 18. As shown in FIG. 27, the Tx PLCP generator 1804 includes a Tx PLCP Demux 2704 and multiple Tx PLCP Modulators 2708. Based on input from the RRC 660, the Tx PLCP Demux 2704 provides data interface streams TxD-j 2716, for j=1 to Jmod, to respective TxPLCP Mod-j 2708 that provide transmit data interface streams Tx-j 2720 to a respective Modulator Core j 1812 (shown in FIG. 18).

In a PTP IBR configuration, such a plurality of data interface streams 2716 and Modulator Cores 1812 may be used, for example, to provide link diversity, such as different sets of RLC parameters, carrier frequencies and/or antenna selections where each set has appropriately multiplexed streams and chains, or to provide probing capability. In a PMP IBR configuration, in addition to the above for PTP, such a plurality of data interface streams and Modulator Cores may also be used at an AE-IBR to optimize data transfer to certain RE-IBRs (or groups of RE-IBRs) by using techniques such as SDMA.

FIG. 28 illustrates an exemplary embodiment of the TxPLCP Mod-j 2708 of FIG. 27. As shown in FIG. 28, the TxPLCP Mod-j 2708 includes a PLCP header generator 2804, a PLCP input buffer 2808, a PLCP controller 2812, a training data library 2816, a padding generator 2820 and a PLCP transmit MUX 2824. In FIG. 28, the PLCP Controller 2812 takes input from the RLC and the RRC as well as the current state of the PLCP Input Buffer 2808 to perform “rate matching” for a given PPDU. Rate matching typically is used to determine a particular choice of modulation and coding scheme (MCS) in view of the instantaneous data transfer demand and recent link conditions. In addition to choosing an MCS, the PLCP Controller 2812 supplies information to the PLCP Header Generator 2804 and directs the PLCP Input

32

Buffer 2808 and the Padding Generator 2820 to supply data to the PLCP Transmit Mux 2824. Depending on the status of the RRC and the instantaneous data transfer demand, the PLCP Controller 2812 also enables training data from the Training Data Library 2816 to be multiplexed into a given PPDU as indicated by certain fields in the PLCP Header. Typically, in some embodiments, the PLCP Header conveys at least the MCS of the PPDU after Block 1, the length of the PPDU payload (for example, either in bytes of payload or in number of blocks plus length in bytes of padding), the MCS of any training data, the selected training dataset, and the number of training data blocks. Optionally, reply acknowledgements (ACKs) or non-acknowledgements (NACKs) for one or more previous receive PPDUs corresponding to a companion receive data interface stream j may also be sent in the PLCP Header. To the extent that the PLCP also concatenates multiple MPDUs and/or fragments thereof, the PLCP Header Generator also includes sufficient information to allow the RxPLCP to recover such MPDUs and/or fragments upon receipt. As described above, FIG. 26 illustrates an exemplary PPDU having a PLCP Header, a PPDU payload including one or more MPDUs or fragments thereof, and a PAD to cause the PPDU to occupy an integer number of Blocks as shown. Not shown in FIG. 26 is a PPDU extension with one or more Blocks of data from the Training Data Library 2816 of FIG. 28 as may be indicated by the PLCP Header.

Note that for Modulator Cores 1812 that divide each transmit data interface stream Tx-j into two or more transmit symbol streams, the PLCP Header of FIG. 28 may be an amalgam of two or more PLCP Headers for respective transmit symbol streams especially in the case where a separate MCS applies to each stream (the “vertical encoding” case shown for example in FIG. 21). Alternatively, a single PLCP Header may be used per Modulator Core j 1812 with parts of the PLCP Header apportioned amongst two or more transmit symbol streams (the “horizontal encoding” case shown for example in FIG. 19). Another alternative would have a single PLCP Header encoded and replicated across all transmit symbol streams of a particular Modulator Core j 1812. In some embodiments, each PLCP Header further includes an FCS (not shown in FIG. 28) to enable the Tx PLCP to determine if the PLCP Header has been demodulated and decoded without error. Note also that the PLCP Controller 2812 may be used in certain embodiments to provide or control the provision of timing signals to the IBR Channel MUX 628 (not shown in FIG. 23) that enable insertion of either or both of the optional Preamble or of the block of Training Pilots as shown in FIG. 26. In exemplary embodiments, a block of Training Pilots, used, for example, to update transmit and/or receive weights by the IBR at an opposite end of a link, may be inserted at the beginning of every transmit superframe, the beginning of every PPDU (as shown in FIG. 26), or at such intervals within a PPDU or superframe as may be communicated to the Tx PLCP generator 1804 or IBR Channel MUX 628 via the RRC 660.

FIG. 29 illustrates an exemplary embodiment of the Rx PLCP analyzer 1808 of FIG. 18. As shown in FIG. 29, the Rx PLCP analyzer 1808 includes a Rx PLCP MUX 2904 and multiple Rx PLCP Demod-j 2908. In FIG. 29, multiple receive data interface streams Rx-j each from respective Demodulator Core j 1812 (see FIG. 18) are input to a respective Rx PLCP Demod-j 2908. After removing the respective PLCP Headers, Padding (if any), and Training Data (if any), each RxD-j is then multiplexed in the Rx PLCP Mux 2904 and transferred to the IBR MAC 612 as Rx Data 620.

FIG. 30 illustrates an exemplary embodiment of the Rx PLCP Demod-j 2908 of FIG. 29. As shown in FIG. 30, the

RxPLCP Demod-j **2908** includes a PLCP header analyzer **3004**, a PLCP output buffer **3008**, a PLCP controller **3012**, a training data analyzer **3016**, and a PLCP receive Demux **3024**. The PLCP Controller **3012** may be shared in some embodiments with the PLCP Controller **2812** of FIG. **28**. The Tx or Rx PLCP Controllers **2812**, **3012** may also be shared across some or all data interface streams. When a PPDU is received from a receive data interface stream Rx-j, the PLCP Header Analyzer **3004** filters certain fields from the PLCP Header and sends this information to the PLCP Controller **3012** (or alternatively, is a part of the PLCP Controller **3012**). This enables the exemplary PLCP Controller **3012** to signal the MCS for subsequent symbols or blocks to Demodulator Core j **1812** as shown in FIG. **30** and to control the PLCP Receive Demux **3024** such that the receive data interface stream Rx-j is directed to the PLCP Output Buffer **3008** or Training Data Analyzer **3016** or ignored in the case of Padding. Further information from the PLCP Header Analyzer **3004** and/or PLCP Controller **3012** enables the Training Data Analyzer **3016** to determine the match between transmitted and received/demodulated/decoded training data and communicate certain metrics such as the number of bit errors and the distribution of errors to the RRC **660**. In embodiments where the PLCP header carries an FCS, the PLCP Header Analyzer **3004** also checks the FCS (not shown in FIG. **30**) and if a failure occurs, signals at least the PLCP Controller **3012**, the RRC **660**, and the RLC **656**. The PLCP Controller **3012** may optionally also inhibit sending Rx-j for the instant PPDU in the case of such an FCS failure or it may send it on with an indicator (not shown) to the IBR MAC that such an FCS failure has occurred.

With reference back to FIGS. **6** and **7**, other IBR elements include the IBR MAC **612**, the Radio Link Control (RLC) **656**, the Radio Resource Control (RRC) **660** and, specific to FIG. **7**, the IBMS Agent **700**. Although IBR embodiments are possible wherein the MAC **612**, RLC **656**, RRC **660** and IBMS **700** are distinct structural entities, more commonly IBRs are realized wherein the IBR MAC **612**, RLC **656**, RRC **660**, IBMS **700** and portions of the IBR Interface Bridge **608** are software modules executing on one or more microprocessors. Note also that in some IBR embodiments that use of a "Software Defined Radio" (SDR) as the IBR Modem **624** and/or IBR Channel MUX **628** or portions thereof may also be realized in software executing on one or more microprocessors. Typically in SDR embodiments, the one or more microprocessors used for elements of the PHY layer are physically separate from those used for the MAC **612** or other layers and are physically connected or connectable to certain hardware cores such as FFTs, Viterbi decoders, DFES, etc.

FIGS. **31** and **32** illustrate exemplary views of the IBRs of FIGS. **6** and **7**, respectively, showing exemplary communications protocols stacks. It will be appreciated that some embodiments have the one or more 802.1 MAC instances, the IBR LLC, the IBR MAC, the RRC, the RLC, the IBMS and all other upper layer protocols implemented as software modules or processes executing on one or more microprocessors within the IBR.

In accordance with the previous description of the IBR Interface Bridge **608**, the 802.1 MAC instances of FIGS. **31** and **32** correspond respectively to Ethernet interfaces **604** and are substantially conventional. Each 802.1 MAC instance is paired with a respective 802.3 PHY instance of which exemplary types include 100 Base-T, 1000 Base-T and various members of the 1000 Base-X family of fiber optic interfaces.

In many exemplary embodiments of the IBR, the IBR Logical Link Control (IBR LLC) layer of FIGS. **31** and **32** is compatible with IEEE 802.2. In accordance with the previous

description of the IBR Interface Bridge, the IBR LLC layer also uses substantially conventional processes to bridge, for example, 802.2 compatible frames across disparate media types as exemplified by the 802.1, the 802.11 and the IBR MAC layers of FIGS. **31** and **32**. Furthermore, in accordance with the previous description of the IBR Interface Bridge, the IBR LLC layer also uses substantially conventional processes to switch and/or load balance 802.2 compatible frames amongst multiple 802.1 MAC instances if available.

FIGS. **31** and **32** depict an exemplary IEEE 802.11 compatible radio (also known as "WiFi") including 802.11 MAC and 802.11 PHY layers and one or more antennas. In some IBR embodiments, such an 802.11 radio, which may be compatible with 802.11b, 802.11a, 802.11g, 802.11n or other 802.11 PHY layer variants, within the IBR may be configured as an access point (AP) to allow significant wireless local area network (WLAN) traffic to be backhauled by the IBR over its wireline or wireless interfaces as appropriate. In some embodiments, the same antennas may be utilized in both the IBR PHY, which typically includes the IBR Modem **624**, IBR Channel MUX **628**, IBR RF **632** and IBR Antenna Array **648** in reference to FIGS. **6** and **7**, and the 802.11 PHY (e.g., via the IBR RF Switch Fabric of FIG. **10** (not shown)). Note further that multiple instances of the 802.11 radio can also be connected to the IBR LLC in an analogous manner to connecting multiple Ethernet interfaces.

Another IBR embodiment may provide a more limited 802.11 radio capability for local configuration purposes (to the exclusion of WLAN traffic and general access). Traditionally PTP and PMP systems have provided a "console" input for local configuration via a command line interface of radios in the field particularly for situations where network configuration is either unavailable or undesirable. However, when an IBR is deployed, for example, on a street light, traffic light, building side, or rooftop, a wired console access with a terminal may be extremely inconvenient and/or dangerous. Thus, in some embodiments, 802.11 radios are deployed as an access point such that terminals including smartphones, tablets, laptops or other portable computing devices with 802.11 station capability can connect to such IBRs in such a console mode. In one embodiment, this deployment is used solely for configuration purposes. Configuration by a terminal with 802.11 station capability is also possible in the case of one or more 802.11 APs deployed for WLAN access purposes. Note further that for the configuration-only 802.11 AP in an IBR, exemplary embodiments may also deploy the 802.11 radio as an Ethernet device on one of the P Ethernet interfaces depicted in FIGS. **6**, **7**, **31** and **32**, or via some other IBR-internal wired or bus interface. In such configuration-only 802.11 AP IBR embodiments, one or more separate antennas for 802.11 use would typically be provided by the IBR. Furthermore, the configuration-only 802.11 interface is not restricted to AP mode only but can also include peer to peer ("IBSS"), AP mode at the terminal, or WiFi direct.

With reference to FIGS. **31** and **32**, typical applications protocols such as Hyper Text Transfer Protocol (HTTP), Simple Network Management Protocol (SNMP), and others can also be implemented using substantially conventional software processes or modules executing on one or more microprocessors within exemplary IBRs. Such applications protocols can get access to or be accessed from any of the depicted network interfaces in FIGS. **31** and **32** via industry-standard Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) transport layer protocols and the Internet Protocol (IP) network layer protocol. TCP, UDP and IP can be implemented using substantially conventional software processes or modules executing on one or more micro-

35

processors within exemplary IBRs. Note that in certain deployments, such as backhaul within a cellular operator's radio access network (RAN), messages to or from HTTP, SNMP or other applications protocols may need to be sent either via other transport and network protocols or via encapsulation of TCP/UDP and IP segments and packets within such other transport and network protocols (not shown in FIGS. 31 and 32). Note further that messages to or from applications protocols in the IBR may also bypass the TCP/UDP and IP protocols, or their equivalents, entirely during a console mode (not shown) using a locally connected terminal.

FIGS. 31 and 32 also illustrate an exemplary element entitled "IBR Control" which is typically implemented as a distribution plane of interconnects, buses and/or controllers that arbitrate direct communications between the RRC and RLC entities with various elements of the IBR MAC and IBR PHY. Because the IBR Control does not pass information amongst such entities by using the communications protocols stack, such information can be communicated with minimal latency for real time control.

FIG. 33 illustrates an exemplary implementation of the IBR MAC 612. The IBR MAC 612 is shown in terms of various functional elements and some of their relationships to each other as well as to the IBR PHY, the IBR LLC, the RRC, the RLC and the IBR Control. Operation of the IBR MAC of FIG. 33 may depend on the type of "superframe" timing utilized. Typically, IBRs send a single PPDU per Modulator Core-j (or per transmit stream k), which may include Training Data blocks as described above, in a given transmit superframe (TxSF) and conversely will receive a single PPDU per Demodulator Core-j (or per receive stream l) in a given receive superframe (RxSF).

As shown in FIG. 33, the IBR MAC 612 includes a IBR MAC management entity 3304, a MAC Tx Buffer and Scheduler 3308, an IBR MAC Control Entity 3312, a MAC Rx Buffer 3316, a decryption block 3320, an encryption block 3324, a MPDU header generator 3328, a first FCS generator 3332, a MPDU header analyzer 3336, a second FCS generator 3340 and a FCS Pass? analyzer 3344.

FIG. 34 illustrates channel activity for a modulator and demodulator pair j with data interface streams Tx-j and Rx-j respectively versus time for an exemplary IBR using FDD with fixed superframe timing. Similarly, FIG. 35 illustrates analogous channel activity for TDD with fixed superframe timing and FIG. 36 illustrates analogous channel activity for TDD and Collision Sense Multiple Access (CSMA) with variable superframe timing.

With reference back to FIG. 33, the operation of the IBR MAC is first described in the context of an example based on TDD/CSMA with variable superframe timing. As shown in FIG. 36, in CSMA transmissions by either the AE-IBR or RE-IBR paired in an active link, whether PTP or PMP, defers to other channel activity detected within its present receive antenna pattern. In practice, this can be achieved by using the receive portion of the IBR PHY to determine in channel and in-view signal energy which if above a threshold causes an inhibit control signal at either a Tx PLCP Mod-j of FIG. 27 (not shown) corresponding to the affected link or at the MAC Tx Buffer and Scheduler of FIG. 33 (not shown) identifying the affected link. For purposes of this description, the latter option is considered. Furthermore, it is assumed that for this specific example every MSDU or frame request at the MAC Tx Buffer and Scheduler results in a single MPDU based on such MSDU or frame request (see, for example, FIG. 4) that is then matched to a single PPDU in the Tx PLCP corresponding to a given link. Other variants are possible with the IBR MAC and IBR PHY as discussed subsequently.

36

With reference again to FIG. 33, the primary source of MSDUs to the MAC Tx Buffer and Scheduler 3308 is typically in the data plane from the IBR LLC 3348. Either MSDUs and/or frame requests in the control plane can originate from the RRC 660 (or indirectly from the optional IBMS Agent 700), the RLC 656, the IBR MAC Management Entity 3304, or the IBR MAC Control Entity 3312. Exemplary details of frame requests and/or MSDUs originating from the RRC 660 or RLC 656 are discussed subsequently.

Exemplary frames originating at the IBR MAC Management Entity 3304 include those associated with management processes such as Association, Authentication and Synchronization. In some cases, the IBR MAC Management Entity 3304 may send an MSDU payload to the MAC Tx Buffer and Scheduler 3308 with a specific frame request type. In other cases, the IBR MAC Management Entity 3304 may send only the specific frame request type wherein all relevant information to be conveyed to the receiving IBR(s) will be present in the MPDU Header to be generated based on the details of the frame request type and other information presently known to the IBR MAC.

In some embodiments, Association and Authentication processes can occur via exchange of management frames in a substantially conventional fashion. A particular RE-IBR may choose to associate with a given AE-IBR by sending an Association Request management frame directed to the AE-IBR based on either advertised information received from such AE-IBR and/or configuration information currently present in the RE-IBR. Upon receipt of an Association Request, the AE-IBR can proceed according to its presently configured policies to associate with or deny association to the RE-IBR. If associated, the AE and RE would exchange authentication frames in substantially conventional fashion to ensure that compatible encryption/decryption keys are present for subsequent frame exchanges.

In exemplary IBRs, an RE-IBR can, for example, associate with a different AE-IBR if its present AE-IBR (or its wireline link interface) fails, the link throughput falls below a minimum threshold, the present AE-IBR forces a disassociation, the present AE-IBR inhibits link resource allocations to the RE-IBR below a minimum threshold, or the RE-IBR becomes aware of a preferred AE-IBR, all as set by certain configuration policies at the time as may be set by the optional IBMS Agent 700 as shown in FIG. 33. Furthermore, certain exemplary IBRs with a plurality of modulator cores can as an RE-IBR maintain a plurality of current associations with multiple AE-IBRs per configuration policies to enable enhanced throughput and/or link diversity.

Another set of exemplary management frames issued by the IBR MAC Management Entity 3304 concerns synchronization, status and presence. Periodically, (as configured or directed by the optional IBMS Agent 700) an exemplary AE-IBR may choose to send a Synchronization Frame that advertises certain capabilities of the AE-IBR including wireline link failure conditions and provides a time stamp reference usable by exemplary RE-IBRs for timing synchronization as either broadcast uniformly across the full directionality possible for the IBR Antenna Array and/or across all current links. Advantageously, particularly for an AE-IBR with multiple associated RE-IBRs in a PMP configuration, such a Synchronization Frame (or other such management frame) can direct one or more RE-IBRs to make internal reference timing offsets such that the time of arrival of transmissions from such RE-IBRs is more optimally aligned for best simultaneous reception at the AE-IBR in either FDD or TDD with fixed superframe timing (see FIGS. 34 and 35). For example, the AE-IBR may determine such

timing offsets within the Acquisition/Synchronization element of the IBR Channel MUX **628** by analyzing preamble samples corresponding to particular RE-IBR transmissions.

With reference again to FIG. **33**, exemplary frames originating at the IBR MAC Control Entity **3312** include those associated with control processes such as Acknowledgement and Access Control. Analogously to the IBR MAC Management Entity **3304**, the IBR MAC Control Entity **3312** may send an MSDU payload to the MAC Tx Buffer and Scheduler **3308** and/or a specific frame request type wherein relevant information can be conveyed to the receiving IBR in the MPDU Header.

In exemplary IBRs, Acknowledgement Frames can provide an ACK or NACK indication for configured frame types recently received. The frames are identified uniquely and then set to ACK or NACK based on the receive path FCS comparison process illustrated in FIG. **33**. To minimize transport overhead, implementation complexity and transmit latency, some embodiments of the IBR utilize a NACK protocol administered by the IBR MAC Control Entity. With reference to FIGS. **34-36**, upon receipt of one or more MPDUs within RxSF(s), where "s" is a superframe sequence number, then if any MPDU has an FCS failure within RxSF(s), a NACK bit is set in the MPDU Header of the first (or each) MPDU within TxSF(s+1). If the originating IBR detects a positive NACK in TxSF(s+1) at its receiver, then such an IBR will re-transmit via its MAC Tx Buffer and Scheduler **3308** the MPDU (or MPDUs) associated with RxSF(s) using link parameters as determined by its RLC **656** at that time. For the FDD case of FIG. **34**, such re-transmitted MPDU (or MPDUs) originally from RxSF(s) would at least partially be received in RxSF(s+2) whereas for the TDD cases of FIGS. **35** and **36**, it would be RxSF(s+1). If there were no FCS failures in RxSF(s), or the originating IBR fails to detect a positive NACK in TxSF(s+1) for any reason, then the originating IBR will clear the MPDU (or MPDUs) associated with RxSF(s) from its MAC Tx Buffer and Scheduler **3308**. Note that unlike conventional ACK protocols, this exemplary IBR NACK protocol does not guarantee 100% delivery of MPDU. However, occasional MPDU delivery failures are typically correctable at a higher layer for data plane MSDUs (i.e., using TCP or an equivalent transport layer protocol), non-catastrophic for management or control plane frames to the IBR MAC Management Entity **3304**, the IBR MAC Control Entity **3312**, the RRC **660**, or the RLC **656**, and preferable to the increased latency and jitter of conventional ACK protocols applied to backhaul for those data plane MSDUs not subject to higher layer correction (i.e. using UDP or an equivalent transport layer protocol). Note further that this NACK protocol is also advantageously used when PLCP Header FCS failures occur by, for example, having the Rx PLCP Demod-j discarding such a failed PPDU but substituting an MPDU with correct link ID but dummy FCS to RxD-j in FIGS. **29** and **30** to trigger an FCS failure at the IBR MAC **612**.

In some embodiments, Access Control Frames are initiated at the IBR MAC Control Entity **3312** to control the behavior of other IBRs with current links to the initiating IBR. Such Access Control Frames can, for example, restrict the rate at which data plane MSDUs are sent, restrict the timing in which data plane MSDUs are sent, or temporarily inhibit further data plane MSDUs from being sent in a local overloading scenario. For example, an exemplary AE-IBR could utilize Access Control Frames to enforce access policies discriminatorily amongst multiple RE-IBRs according to certain configuration parameters such as a Service Level Agreement (SLA). Such access policies may also be set via the optional IBMS Agent **700** as shown in FIG. **33**.

With reference to FIG. **33**, the MAC Tx Buffer and Scheduler **3308** temporarily stores data plane MSDUs from the IBR LLC **3348** as well as frames or frame requests from the IBR MAC Management Entity **3304**, IBR MAC Control Entity **3308**, RRC **660** and/or RLC **656** until such data can be scheduled for transmission within one or more MPDUs. Such scheduling is advantageously performed based on policies either configured locally or optionally communicated from the IBMS Agent **700** via the RRC **660** and the IBR Control **3352**. In some embodiments, data plane MSDUs from the IBR LLC **3348** usually are the lowest priority in the queue for scheduling composition of the next MPDU. If the rate of MSDU delivery from the IBR LLC **3348** exceeds the instant link delivery capability, then to avoid buffer overflow the MAC Tx Buffer and Scheduler **3308** may have an Xon/off control feedback to the IBR LLC **3348** to cause the IBR LLC **3348** to stop sending MSDUs until the appropriate buffer status returns. This may result in the IBR LLC **3348** causing certain other interfaces such as an Ethernet or WiFi interface to reduce MSDU traffic. To the extent that the ratio of control plane traffic to data plane traffic is small, as is by design for some embodiments, then scheduling priority amongst frames other than MSDUs from the IBR LLC **3348** is unimportant and "first-in, first-out" is sufficient.

Note that the foregoing embodiment of the MAC Tx Buffer and Scheduler **3308** performs scheduling based on frame type without regard for other information such as the PCP within an 802.1 MAC frame as may be present in MSDUs from the IBR LLC **3348**. The IBR LLC **3348** in its bridging capacity may forward MSDUs to the MAC Tx Buffer and Scheduler **3308** in order based on PCP (or other such QoS or CoS indicators) so that the MAC Tx Buffer and Scheduler **3308** need not repeat the queuing exercise. In alternative embodiments, such QoS prioritization of MSDUs can also be performed at the MAC Tx Buffer and Scheduler **3308** instead.

Upon scheduling an MPDU for transmission as described above, the MAC Tx Buffer and Scheduler **3308** causes the MPDU Header Generator **3328** to compose an MPDU Header specific to the pending MPDU. In conventional IEEE 802 based communications systems, such an MPDU Header would include at least the physical address (also known as the "MAC address" or the "hardware address") of both the origination and destination IBRs. However, sending such physical addresses (typically 48 or, more recently, 64 bits) in every MPDU unduly burdens the IBR links with unnecessary overhead that reduces MAC efficiency. Thus, in some embodiments, a Link Identifier (LID) is substituted in every MPDU header instead. Exemplary LID implementations can be as few as 16 bits in length. For example, each exemplary LID may include an AE-IBR identifier of 10 bits and an identifier of 6 bits for an RE-IBR presently associated with the AE-IBR. This is possible because in some embodiments the IBRs are configured in view of their fixed geographic positions in the field as set at time of deployment or optionally controlled by the IBMS Agent **700** such that no overlapping AE-IBR identifiers are within radio range of RE-IBRs possibly associated with them. The AE-IBR may assign a locally unique RE-IBR association identifier field as part of the association process. For unassigned links, all zeros (or ones) can be sent for LID and then the frame payload, for example a management frame used in the association process, can include the full physical addresses as appropriate. Note that even if, in the alternative, that longer (possibly 24 bits or 32 bits) "regionally-unique" or even "globally-unique" LIDs were used, then because the overall number of worldwide backhaul links is generally much less the overall number of worldwide network devices,

such extended length LIDs can still be much shorter than traditional IEEE 802 based addressing schemes.

A Frame Type Identifier (FTI) may be placed in the MPDU header by the MPDU Header Generator **3328**. In one embodiment, the FTI is no more than 5 bits total. In one particular embodiment, the FTI is part of a Frame Control Field (FCF) of 8 bits, and the other 3 bits include 1 bit for the NACK protocol control bit indicator (set to 1 if the previous RxSF(s) had an MPDU with an FCS failure), 1 bit to indicate if the instant frame of the MPDU payload is encrypted, and 1 bit to indicate if the instant frame of the MPDU payload is the last frame (LF) in the payload. Alternatively, the 1 bit for the NACK can be 1 bit for the ACK indicator (set to 1 if the previous RxSF(s) had an MPDU without an FCS failure) if an ACK protocol is used for the instant MPDU FTI. Following this FCF byte, a 16 bit Fragment and Length Indicator (FLI) is placed sequentially in the MPDU header, wherein, for example, 3 bits of the FLI indicates by 1 bit if the instant frame payload is the last fragment and by 2 bits the fragment sequence number and 13 bits indicate the instant frame payload length in bytes. Following the 2 FLI bytes, an 8 bit Frame Sequence Number (FSN) is placed sequentially in the MPDU header. The FSN are typically sequentially generated except where repeated for those frame payloads sent as fragments. If LF=1 in the initial FCF byte of the MPDU header (as would be the case for the single MSDU of frame payload per MPDU scenario described above for TDD/CSMA), then the MPDU header is complete. If the MAC Tx Buffer and Scheduler **3308** is configured to permit concatenation of MSDUs or other frame payloads up to some maximum MPDU payload (as would be compatible with many TDD/CSMA deployments), then LF=0 when FLI describes an instant frame payload length sufficiently low to allow another available frame payload to be concatenated within the maximum MPDU payload length and an additional FCF, FLI and FSN combination would be generated at the MPDU Header Generator **3328** and repeated until an FCF with LF=1 is encountered. Note this process of concatenated FCF, FLI and FSN fields within an MPDU header corresponding to concatenated frame payload can also be advantageously applied to fixed superframe timing in either FDD or TDD as illustrated in FIGS. **34** and **35**.

The MAC Tx Buffer and Scheduler **3308** further provides the one or more frame payloads that form the MPDU payload to the Encryption element **3324** to be encrypted on a frame payload by frame payload basis as indicated in the FCF using substantially conventional encryption algorithms. An exemplary algorithm suitable for the IBR is the Advanced Encryption Standard (AES) which has been published by the National Institute of Standards and Technology. In one embodiment, the IBRs use AES with a 256 bit key length. Other key lengths and other encryption algorithms are also possible for exemplary IBRs. Exemplary IBRs can also employ encryption for all frames after encryption keys are exchanged during authentication (and even including association and authentication frames to the extent encryption keys sufficient at least for association and authentication are provided to IBRs via, for example, factory setting or console mode interface).

The encrypted (and/or unencrypted as desired) frame payload(s) and the MPDU header are then concatenated together as shown in FIG. **33** and then passed through an FCS Generator **3332** that generates an FCS (e.g., of at least 32 bits in length). Alternatively, some IBR embodiments may also encrypt (and decrypt) MPDU headers. Other IBR embodiments may attain some measure of MPDU header privacy (and PLCP header privacy) by setting scrambler and descrambler parameters (see FIGS. **19-22**) derived from, for

example, the encryption keys and/or LID. The FCS is then appended following the MPDU header and frame payload(s) as also shown in FIG. **33** to complete the composition of an MPDU to be sent to the IBR PHY **3356** over the Tx Data interface.

Note also that for those frame types, particularly certain management frames originating at the AE-IBR of a PMP configured deployment, that are intended to be broadcast to all current links at an IBR, such frame payloads may be distributed to all such links in parallel at the MAC Tx Buffer and Scheduler **3312** such that the same frame payload is provided to at least one MPDU corresponding to each current link. The IBRs typically generate very little broadcast traffic and most Ethernet or WiFi broadcast traffic on the other IBR interfaces is filtered at the IBR LLC **3348**.

With reference to FIG. **33**, the receive path of the exemplary IBR MAC performs essentially the reverse operations described above in the transmit path. Starting with the IBR PHY **3356**, after the Rx PLCP analyzer **1808** (see FIG. **18**) recovers an MPDU, it is passed via the Rx Data interface in FIG. **33** to a splitter **3360** within the IBR MAC. The splitter **3360** removes the trailing bits of the MPDU corresponding to the FCS, or Rx FCS. The remainder of the MPDU minus Rx FCS is passed through an FCS Generator **3340**, the result of which is compared at FCS Pass? analyzer **3344** to determine if it is identical to Rx FCS. If it is identical, then a known-good (barring the extraordinarily rare false positive FCS event) MPDU is received. The FCS Pass? status is communicated to the IBR MAC Control Entity **3312**, the MAC Rx Buffer **3316**, the MPDU Header Analyzer **3336**, and the RRC **660** and RLC **656** (for analytical purposes) as shown in FIG. **33**. The MPDU Header Analyzer **3336** receives from a second splitter the leading bits of the MPDU corresponding to a minimum length MPDU header described above. If LF=1, then the remaining payload portion of the MPDU is passed to the Decryption element **3320** (or bypassed if appropriate) for decryption when the encryption indicator bit of the FCF is set. The MPDU Header Analyzer **3336** may also verify (or cause to be verified) that the payload length corresponds to that described in the FLI field of the MPDU Header and that the LID is valid. Upon decryption (as appropriate), the MPDU payload is passed to the MAC Rx Buffer **3316** which then forwards the MSDU or frame payload to the IBR MAC Management Entity **3304**, the IBR MAC Control Entity **3312**, the IBR LLC **3348**, the RRC **660** or the RLC **656** as appropriately directed by the MPDU Header Analyzer **3336** based on LID and FTI. The MPDU Header Analyzer **3336** may also directly signal the IBR MAC Control Entity **3312** with the status of the NACK bit in the exemplary MPDU header described above. In the event that the instant frame payload is a fragment, as evident by the FLI field in the exemplary MPDU header described above, then the MPDU Header Analyzer **3336** instructs the MAC Rx Buffer **3316** based on the FSN field to either store the fragment, concatenate the fragment to one or more other stored fragments, or deliver the multiple fragments if the instant fragment has the last fragment indicator within the FLI field as described above. Optionally if a positive ACK protocol is used for 100% delivery of every MPDU (versus the superframe NACK protocol described above), then the MDPU Header Analyzer **3336** also verifies the FSN for duplicated detection and discards (or causes to be discarded) such duplicate payloads.

In the event that the MPDU being received has multiple concatenated MSDUs and/or frame payloads as described optionally above, then the MPDU Header Analyzer **3336** interacts with the second splitter via the feedback signal shown to continue parsing off consecutive MDPU header

41

bytes, for example the repeated exemplary FCF, FLI and FSN fields, until an FCF with LF=1 is encountered. The above described process for the MPDU Header Analyzer 3336 directing the Decryption element 3320 and the MAC Rx Buffer 3316 to deliver the multiple payloads is then repeated either serially or in parallel as desired until all contents of the MPDU payload have been resolved by the IBR MAC.

In the event that the FCS Pass? analyzer 3344 described above determines an FCS failure then the receive path of the exemplary IBR MAC operates differently depending on certain options that may be designed into the IBR or selected based on configuration policies. For the superframe NACK protocol described above, the FCS failure is directly communicated to the RRC 660 and RLC 656 (via the IBR Control 3352 of FIGS. 31-33) for their analytical purposes and is directly communicated to the IBR MAC Control Entity 3312 so that TxSF(s+1) is sent with NACK=1 in the exemplary MPDU header described above. Optionally, the MPDU Header Analyzer 3336 may attempt to determine if an apparently valid MPDU header is present, and, if so, attempt to deliver certain MSDUs or frame payloads as described above for the case of a known-good MPDU. However, as directed by configuration policies, certain frame types may be discarded. For example, in one embodiment, only MSDUs to the IBR LLC 3348 that correspond to apparently valid MPDU header fields would be delivered in an FCS failure situation, and even then only after buffering such MSDUs for one additional superframe period to determine if a duplicate FSN MSDU may be available due to the re-transmission attempt associated with the superframe NACK protocol described above. Alternatively, in the case of an ACK protocol that guarantees 100% delivery of known-good MPDUs, then the entire MPDU payload is discarded.

For the types of data throughput rates and superframe lengths that are practical, the single payload MPDU per superframe example described above for the TDD/CSMA example of FIG. 36 is not an optimum choice for the fixed superframe timing examples for FDD in FIG. 34 or TDD in FIG. 35. One alternative uses concatenated multi-payload frame MPDUs up to some maximum length and then forwards such MPDUs to the Tx PLCP to compose PPDU as described previously. Another alternative, compatible with the superframe NACK protocol described above, would have the MAC Tx Buffer and Scheduler 3308 utilize the same information from the RLC 656 and the RRC 660 available to the Tx PLCP and thus compose a single multi-payload MPDU that will, with appropriate padding added in the Tx PLCP to an integer block count, maximally occupy the available PPDU for the pending superframe.

Note that for the TDD case depicted in FIG. 35 with fixed superframe timing, IBRs are capable of operating in such modes wherein the superframe timing in the forward link from AE to RE may be different than in the reverse link. For a PMP AE-IBR this typically requires that all RE-IBRs adopt the same superframe timing parameters in association with this particular AE-IBR. Some RE-IBRs, in either PTP or PMP configurations, may be simultaneously associated with two or more AE-IBRs to the extent such IBRs permit this policy and have sufficient radio resources to maintain such links. Given that multiple AE-IBRs associated with such RE-IBRs can adopt different TDD fixed superframe timing parameters relative to each other, such timing parameters may be communicated to all RE-IBRs via management frames originating in the IBR MAC Management Entity 3304. Similarly, to the extent an RE-IBR is limited in superframe timing access due to such multiple associations, such an RE-IBR can com-

42

municate such restrictions to the AE-IBRs via management frames from its IBR MAC Management Entity 3304.

For PMP configurations, the AE-IBR can advantageously transmit to or receive from multiple RE-IBRs simultaneously in a given radio channel using SDMA as described above. To the extent that the number of RE-IBRs associated with an AE-IBR exceeds either the AE-IBR's available SDMA resources or certain RE-IBRs are not spatially separable by the AE-IBR, then simultaneous transmissions to such RE-IBRs would require multiple radio channels which in practice is often either undesirable or impractical. Another alternative for maintaining links to RE-IBRs wherein such SDMA capabilities are being exceeded is through the use of Time-Division Multiplexing (TDM). For either an FDD or TDD system with fixed superframe timing, such as depicted in FIGS. 34 and 35, an AE-IBR can use TDM to designate certain "sub-frames" within a given superframe applied to a modulator/demodulator resource pair *j* to particular LIDs. Management frames from the AE-IBR's IBR MAC Management Entity 3304 can, in exemplary embodiments, inform affected RE-IBRs of such LID to TDM sub-frame mappings and associated subframe timing parameters. In some embodiments, any RE-IBR associated with two or more AE-IBRs simultaneously could request that such LID to TDM sub-frame mappings and/or timing parameters account for its local timing preferences for optimally maintaining such simultaneous links. The MAC Tx Buffer and Scheduler 3308 may compose MPDUs in view of TDM sub-frame timing parameters that are assigned such that the TxPLCP can optimally utilize an integer number of transmit blocks in each sub-frame. The number of sub-frames per superframe are typically relatively low (e.g., usually less than 5).

To the extent that IBRs deployed in the field are within co-channel interference range of each other and configured to use overlapping channels in general or periodically, such IBRs advantageously synchronize their TDD fixed superframe timing to minimize simultaneous co-channel TxSF/RxSF overlaps amongst disparate links. Several exemplary synchronization strategies may be used by such IBRs to align superframe timing boundaries in such scenarios. For example, if "free-running" or "self-synchronizing" IBRs are able to detect at least a preamble, a training block, a PLCP header, an unencrypted MPDU header, or other information such as management frames with timing information that may be decipherable without link-specific encryption keys that correspond to links involving a different AE-IBR, then the slower in time IBRs may adopt the superframe timing boundary and cadence of the faster in time IBRs. At the AE-IBR, which may act as a local timing master, this can be performed directly by making a timing offset and communicating it to its associated RE-IBR(s). At the RE-IBR, which may be able to detect disparate link information otherwise undetectable at its associated AE-IBR, the RE-IBR can inform its AE-IBR via a management frame of any timing offset necessary to obtain local disparate AE-IBR co-channel superframe timing synchronization. It will be appreciated that this process is ongoing because after synchronizing, the reference clocks in the AE-IBRs inevitably will drift differently over time.

In certain field deployment scenarios, IBRs located in the same regional area may be capable of undesirably interfering with each other at ranges beyond their ability to detect and synchronize as described above. An alternative synchronization strategy better suited to this situation would utilize a network-wide central synchronization capability. One example of this would be the use of Global Positioning Satellite (GPS) timing at each AE-IBR. GPS is more than

adequate in terms of timing accuracy for the needs of synchronizing superframe timing boundaries. However, GPS adds cost, size, power consumption and form factor constraints that may be undesirable or unacceptable to some IBRs. Note further that because IBRs are designed to operate in street level deployments with obstructed LOS, GPS may fail to operate in places where RE-IBRs function normally. Another alternative would be to use a system-wide synchronization technique such as SynchE or IEEE 1588v2. In this scenario, AE-IBRs are configured to derive timing parameters in a consistent fashion. Alternatively, the AE-IBRS include IBMS Agents capable of coordinating such configurations when co-channel operation in a mutual interfering deployment is encountered.

In the deployment scenario where multiple AE-IBRs are co-located (e.g., at a single pole, tower, building side or rooftop), even if such IBRs are configured to avoid co-channel operation, at least some form of local superframe timing synchronization for TDD may be utilized to avoid overloading receiver RF/analog circuits from simultaneous in-band transmissions and receptions amongst the co-located IBRs. One exemplary strategy for the above synchronization would be to distribute by hard wiring a local superframe timing synchronization signal which can be configured at or arbitrarily assigned to any one of the co-located IBRs.

Note that the foregoing descriptions and figures for exemplary IBR embodiments have provided minimal internal details on the distribution of various clocks and timing references amongst the structural and functional elements of the IBR. These exemplary embodiments can all be realized using substantially conventional strategies for internal timing distribution.

As described above, IBRs with fixed superframe timing may use a NACK protocol wherein a previous superframe FCS failure in receive causes NACK=1 in an MPDU header of the respective link in transmitting its next sequential superframe PPDU back to the sender of the MPDU received in error. If the original sender detects NACK=1, then a re-transmission of the previous superframe PPDU contents occurs at the direction of its MAC Tx Buffer and Scheduler 3308. Otherwise, the original sender discards the previous superframe PPDU contents at its MAC Tx Buffer and Scheduler 3308. This approach is different from many conventional wireless data networking protocols that use ACK receipt to guarantee 100% delivery at the MAC layer even if theoretically unbounded re-try attempts may be required. This fixed superframe NACK protocol is similar to using an ACK protocol with a "time to live" of effectively one superframe duration after initial transmission. This approach advantageously bounds the latency at a very low length for processing frames through the IBR MAC without resorting to reliability of simply the raw frame error rate. By allowing one immediate re-transmission opportunity, this fixed superframe NACK protocol effectively produces a net frame error rate that is the square of the raw frame error rate. For example, if the raw frame error rate were 10^{-3} (1 in 1000), the net frame error rate per the fixed superframe NACK protocol should be approximately 10^{-6} (1 in 1,000,000).

To improve the reliability of the fixed superframe timing NACK protocol even further, IBRs may have the TxPLCP set a 1 bit field in the PLCP header(s) for the Modulator Core j corresponding to the LID with NACK=1 (or with a previous RxSF(s) PLCP header FCS failure as described above). This approach advantageously exploits the fact that PLCP headers, typically sent at the most reliable MCS, are of short duration and always sent immediately after the Training Block 0 (see FIG. 26) where the channel equalization coefficients most

accurately reflect current conditions. Thus, a NACK bit sent in a PLCP header is likely to be received accurately even if the PPDU payload that includes the MPDU header NACK field is not. In the case where NACK=1 in a PLCP header, an exemplary PLCP Header Analyzer 3004 (see FIG. 30) will signal the exemplary MAC Tx Buffer and Scheduler 3308 via the IBR Control 3352. In the case where an originator of an MPDU does not receive a valid NACK field in either the PLCP header or the next MPDU header of a subsequent transmit superframe, then the MAC Tx Buffer and Scheduler 3308 may choose to re-send the pending MPDU(s), possibly at a more reliable MCS as directed by the RLC 656, or may choose to discard the pending MPDU(s) so as not to cascade latency for further frames awaiting transmission. Such a decision may be made based on configurable policies that may vary with current conditions or be updated by the IBMS Agent 700. To the extent that a "blind" retransmission is made due to an invalid NACK field wherein the actual NACK field value was NACK=0, then the receive path MPDU Header Analyzer 3336 at the destination IBR will discard the re-sent MPDU(s) based on duplicate detection of the FSN at the MPDU Header Analyzer 3336 shown in FIG. 33. The MPDU Header Analyzer 3336 may also report incidents of duplicate detections of MPDUs for a given LID to the RRC 660 and RLC 656 as shown in FIG. 33. This provides an advantageous additional local indication to the RRC 660 and RLC 656 of problems being encountered in the transmit direction from such an IBR for the particular LID.

With reference to FIGS. 31 and 32, the RRC 660 and RLC 656 interact with the IBR MAC 612 and various elements of the IBR PHY both via "normal" frame transfers and direct control signals via the conceptual IBR Control plane. Both the RRC 660 and the RLC 656 may execute concurrent control loops with the respective goals of optimizing radio resource allocations and optimizing radio link parameters for current resources in view of the dynamic propagation environment conditions, IBR loading, and possibly system-wide performance goals (via the optional IBMS Agent 700). It is instructive to view the RLC 656 as an "inner loop" optimizing performance to current policies and radio resource allocations for each active link and to view the RRC 660 as an "outer loop" determining if different policies or radio resource allocations are desirable to meet overall performance goals for all IBRs currently interacting with each other (intentionally or otherwise). Typically both the RRC 660 and the RLC 656 are implemented as software modules executing on one or more processors.

The primary responsibility of the RLC 656 in exemplary IBRs is to set or cause to be set the current transmit MCS and output power for each active link. In one exemplary embodiment described above, the RLC 656 provides information to the TxPLCP that enables, for example, a PLCP Controller 2812 in FIG. 28 to set the MCS for a particular LID. In another embodiment compatible with the single MPDU for fixed superframe timing with NACK protocol example of IBR MAC operation, the RLC 656 determines the MCS for each LID and then communicates it directly to the MAC Tx Buffer and Scheduler 3308 of FIG. 33 and the PLCP Controller 2812 for every TxPLCP Mod-j of FIGS. 27 and 28. The RLC 656 causes the transmit power of the IBR to be controlled both in a relative sense amongst active links, particularly of interest for the AE-IBR in a PMP configuration, and also in an overall sense across all transmits chains and antennas. In an exemplary IBR embodiment, the RLC 656 can cause such transmit power control (TPC) as described above by setting parameters at the Channel Equalizer Coefficients Generator 2332 of FIG. 23 (for relative power of different simultaneous modu-

lation streams), at each active transmit RF chain Tx-m **636** of FIGS. **6**, **7** and **16** (for relative power of different simultaneous RF chains) and at each Front-end PA **1104**, **1204** within the IBR Antenna Array of FIGS. **6**, **7**, **10-12** (for total power from all antennas).

In some embodiments, the RLC **656** can determine its MCS and TPC selections across active links based on information from various sources within the IBR. For example, the IBR MAC can deliver RLC control frames from other IBRs with information from such other IBRs (for example, RSSI, decoder metrics, FCS failure rates, etc.) that is useful in setting MCS and TPC at the transmitting IBR. Additionally, such RLC control frames from an associated IBR may directly request or demand that the RLC in the instant IBR change its MCS and/or TPC values for transmit directly on either a relative or absolute basis. For TDD IBR deployments, symmetry of the propagation environment (in the absence of interference from devices other than associated IBRs) makes receiver information useful not only for sending RLC control frames to the transmitting IBR but also for use within the receiving IBR to set its transmitter MCS and/or TPC. For example, the FCS Pass? analyzer **3344** and MPDU Header Analyzer **3336** of the exemplary IBR MAC in FIG. **33** can supply the RLC **656** with useful information such as FCS failures, duplicate detections and ACK or NACK field values. Similarly, each PLCP Header Analyzer **3004** as shown in FIG. **30** can send the RLC FCS failure status on the PLCP Header and/or PPDU payload(s). Another possibility is to have the PLCP Header carry a closed-loop TPC control field requesting a TPC step up or down or stay the same. Additionally, the Channel Equalizer Coefficients Generator **2332** in computing channel weights can provide SNR and/or SINR information per demodulator stream that can help the RLC set MCS and/or TPC as shown in FIG. **23**. Each decoder within each demodulator core *j* can provide the RLC **656** with decoder metrics useful for MCS and/or TPC as indicated by FIGS. **20** and **22** for example. And each receive chain Rx-n **640** of FIG. **17** can provide receive signal strength indicator (RSSI) values helpful especially for determining TPC requests back to the transmitting IBR.

The actual MCS values are typically selected from a finite length list of modulation types and coding rates. Exemplary IBRs can use QAM ranging from 2-QAM (better known as BPSK), through 4-QAM (better known as QPSK), 16-QAM, 64-QAM, 256-QAM and 1024-QAM. Exemplary IBRs can use a base coding rate of 1/3 or 1/2 and then can use "puncturing" (wherein predetermined bit positions are deleted in transmit and replaced by dummy bits in receive) to derive a set of effective coding rates of, for example only, 1/2, 2/3, 3/4, 5/6, 7/8, and 9/10. In typical embodiments, the lowest MCS index corresponds to the lowest available QAM constellation size and the lowest available coding rate (i.e. the most reliable transmission mode) and the highest MCS index corresponds to the converse.

The TPC absolute range tends to be lower for IBRs than that desired for many conventional wireless networking systems operating in obstructed LOS due to the more limited range of separations between AE-IBRs and RE-IBRs for backhaul applications (i.e. backhaul radios are almost never placed in close proximity to each other). The relative variation in desired power between active links at an AE-IBR may also be limited in range particularly by the transmit DACs.

Many possible algorithms are known for generally relating information provided to the RLC **656** as described above to selecting MCS and TPC values. In dynamic propagation environments, averaging or dampening effects between channel quality information and MCS changes are advantageously

utilized to avoid unnecessarily frequent shifts in MCS. To the extent that an IBR is operating at below the maximum allowable TPC value, it is generally advantageous to permit TPC to vary more quickly from superframe to superframe than the MCS. However, if operating at maximum allowable TPC, then it is often advisable to immediately down select MCS to a more reliable setting upon detection of an MPDU FCS failure and/or a NACK=1 condition. Conversely, up selecting MCS is usually performed only after repeated superframe metrics indicating a high likelihood of supporting an MCS index increase. At the limit, where RLC **656** has reached maximum TPC and minimum MCS (most reliable mode), to maintain ongoing link reliability, the imperative increases for the RRC **660** to allocate different resources to enable the RLC **656** to operate again in MCS and TPC ranges with margin for temporal channel impairment.

The primary responsibility of the RRC **660** is to set or cause to be set at least the one or more active RF carrier frequencies, the one or more active channel bandwidths, the choice of transmit and receive channel equalization and multiplexing strategies, the configuration and assignment of one or more modulated streams amongst one of more modulator cores, the number of active transmit and receive RF chains, and the selection of certain antenna elements and their mappings to the various RF chains. Optionally, the RRC may also set or cause to be set the superframe timing, the cyclic prefix length, and/or the criteria by which blocks of Training Pilots are inserted. The RRC **660** allocates portions of the IBR operational resources, including time multiplexing of currently selected resources, to the task of testing certain links between an AE-IBR and one or more RE-IBRs. The RRC **660** evaluates such tests by monitoring at least the same link quality metrics as used by the RLC **656** as evident in the exemplary embodiments depicted in FIGS. **6**, **7**, **17**, **18**, **20**, **22**, **23**, **29**, **30**, **31**, **32** and **33**. Additionally, in some embodiments, additional RRC-specific link testing metrics such as those produced by the exemplary Training Data Analyzer described above for FIG. **30** in relation to the exemplary Training Data Library of FIG. **28** are used. The RRC **660** can also exchange control frames with a peer RRC at the other end of an instant link to, for example, provide certain link testing metrics or request or direct the peer RRC to obtain link specific testing metrics at the other end of the instant link for communication back to RRC **660**.

In some embodiments, the RRC **660** causes changes to current resource assignments in response to tested alternatives based on policies that are configured in the IBR and/or set by the optional IBMS Agent **700** as depicted for example in FIGS. **7**, **32**, and **33**. An exemplary policy includes selecting resources based on link quality metrics predicted to allow the highest throughput MCS settings at lowest TPC value. Additional exemplary policies may factor in minimizing interference by the instant link to other AE-IBR to RE-IBR links (or other radio channel users such as conventional PTP radios) either detected at the instant IBRs or known to exist at certain physical locations nearby as set in configuration tables or communicated by the optional IBMS Agent **700**. Such policies may also be weighted proportionately to reach a blended optimum choice amongst policy goals or ranked sequentially in importance. Also as discussed above regarding the RLC **656**, the RRC **660** may have policies regarding the fraction of IBR resources to be used for RRC test purposes (as opposed to actual IBR backhaul operations) that factor the current RLC settings or trajectory of settings.

The RRC **660** in one IBR can communicate with the RRC in its counterpart IBR for a particular link AE-IBR to RE-IBR combination. For example, the RRC **660** sends RRC control

frames as discussed for the exemplary IBR MAC of FIG. 33 or invokes a training mode within the exemplary Tx PLCP as discussed for FIG. 28.

In some embodiments, for either PTP or PMP deployment configurations, the selection of either the one or more active RF carrier frequencies used by the RF chains of the IBR RF, the one or more active channel bandwidths used by the IBR MAC, IBR Modem, IBR Channel MUX and IBR RF, the superframe timing, the cyclic prefix length, or the insertion policy for blocks of Training Pilots is determined at the AE-IBR for any given link. The RE-IBR in such an arrangement can request, for example, an RF carrier frequency or channel bandwidth change by the AE-IBR by sending an RRC control frame in response to current link conditions at the RE-IBR and its current RRC policies. Whether in response to such a request from the RE-IBR or due to its own view of current link conditions and its own RRC policies, an AE-IBR sends the affected RE-IBRs an RRC control frame specifying at least the parameters for the new RF frequency and/or channel bandwidth of the affected links as well as a proposed time, such as a certain superframe sequence index, at which the change-over will occur (or alternatively, denies the request). The AE-IBR then makes the specified change after receiving confirmation RRC control frames from the affected RE-IBRs or sends a cancellation RRC control frame if such confirmations are not received before the scheduled change. In some deployment situations, the RRC policy condition causing the change in the RF carrier frequency and/or channel bandwidth for a particular LID may be a directive from the IBMS Agent 700.

The selection of other enumerated resources listed above at an IBR can generally be made at any time by any given IBR of a link. Additionally, requests from the opposite IBR can also be made at any time via RRC control frames. An RE-IBR typically attempts to utilize all available modulator and demodulator cores and streams as well as all available RF chains to maximize the robustness of its link to a particular AE-IBR. In an RE-IBR embodiment where at least some redundancy in antenna elements amongst space, directionality, orientation, polarization and/or RF chain mapping is desirable, the primary local RRC decision is then to choose amongst these various antenna options. An exemplary strategy for selecting antenna options (or other enumerated resources listed previously) at the RE-IBR is to apply alternative selections of such resources to the Training Data portion of PPDU's described above in relation to the Tx PLCP and FIGS. 27 and 28. The RRC 660 at the RE-IBR then compares link quality metrics such as those used by the RLC 656 and training data metrics from the exemplary Training Data Analyzer of FIG. 30 to determine if the alternatively-selected resources are likely to result in improved performance based on current IBR RRC policies compared to the known performance of the currently-selected resources of the instant link. If the RRC 660, based on one or more such alternative selection tests, decides that the instant link's performance goals are likely to improve by using the alternately-selected resources, then the RRC 660 at such RE-IBR can cause such selections to become the new current settings at any time and without requiring notification to the AE-IBR.

For the RE-IBR alternate resource selection process described above applied to a TDD configuration, channel propagation symmetry for a given link (if interference is ignored) may make changing to a corresponding set of resources for transmit from such RE-IBR as have been alternatively-selected for receive preferable. However, this is generally not true for an FDD configuration or a scenario where unknown interference represents a significant channel

impairment or where a PMP AE-IBR has simultaneous links to other RE-IBRs. In such scenarios, an RE-IBR may notify the AE-IBR when such RE-IBR is testing alternately-selected resources in a portion of its Tx PPDU's, whether in response to an RRC control frame request by the AE-IBR or by such RE-IBR's own initiative, and then receive a return RRC control frame from the AE-IBR that either reports the measured link quality metrics observed at the AE-IBR and/or directs the RE-IBR to adopt the alternatively-selected resources for future Tx PPDU's from the RE-IBR on such link.

For the PTP configuration, an AE-IBR performs its RRC-directed alternate resource selections using substantially the same processes described above for the RE-IBR but with the roles reversed appropriately. In the PMP configuration, an AE-IBR may utilize similar RRC-directed testing of alternate resource selections across its multiple current links but to the extent that such links depend concurrently on certain resources, the decision to actually change to different resources may be based on policies applied to the benefit of all current links. One strategy for PMP operation of IBRs is to use the maximum possible RF chain and antenna element resources at all times at the AE-IBR and then optimize selectable resources at the RE-IBRs to best achieve RRC policy goals.

Note that in some deployment situations, spectrum regulations, such as those set by the Federal Communications Commission (FCC) in the USA, may require active detection of and avoidance of interference with other users of the spectrum (such as, for example, radar systems). The process of detecting such other co-channel spectrum users and then changing RF carrier frequencies to another channel void of such other uses is commonly called Dynamic Frequency Selection (DFS). Spectrum regulations may require that a DFS capability operate in a certain manner in response to certain interference "signatures" that may be detected at the receiver of a certified radio for such spectrum. For example, some regulations require that upon detection of certain pulse lengths and received powers that certified radios change immediately to another channel known from some minimum observation time not to have such interferers in operation. In some exemplary IBR implementations, such observations of alternative channels in advance of needing to make a change can be performed by time division multiplexing certain RF chain and antenna resources to make such measurements using RSSI and/or channel equalization metrics reported to the RRC 660. In some embodiments, the AE-IBR and the one of more associated RE-IBRs coordinate such observations using RRC control frames to minimally disrupt backhaul operations and maximally increase the aggregate observation time and improve the observation accuracy. In other exemplary IBR embodiments, at least one IBR, typically the AE-IBR, has at least one Rx-n chain, one antenna and possibly one demodulator core path through the IBR Channel MUX dedicated to such spectral observations.

In embodiments with the optional IBMS Agent 700, the above channel observation techniques of the RRC 660 can also be used in a "probe in space" mode of operation, either at one IBR or coordinated amongst multiple IBRs, to observe and record RF channel activity within designated portions of the addressable bands of operation. Such spectral analysis information may be passed from the RRC 660 to the IBMS Agent 700 for further analysis and possibly communication to other IBMS Agents or to other databases or expert systems elsewhere within the IBR operator's private network or to an external network such as the Internet.

Note also that the DFS operation described above is desirable for exemplary IBRs operating in spectrum bands that do

not require DFS explicitly. For example, if IBRs are deployed in licensed spectrum where conventional PTP links operate, such conventional links generally lack the RF carrier frequency agility and radio resource control intelligence of the IBR. Even if the interference immunity capabilities of the IBR through the advantageous combinations of elements described herein is sufficient to reject the interference caused by such conventional PTP links at the IBR receiver, it is still desirable to have the RRC 660 perform DFS to avoid the converse scenario where the IBRs are interfering with the conventional PTP links. This may be advantageous because it minimizes licensed band user conflicts especially if a different operator uses the conventional PTP equipment from that operating the IBRs. The presence of the conventional PTP link may be detected in normal operation of the one or more IBRs using a particular RF carrier frequency channel or may be communicated to the RRC 660 via the optional IBMS Agent 700 that has gathered the information from another source. An exemplary technique that the RRC 660 can use in an IBR where $N > L$ and the instant SINR is approximately the same as the SNR (i.e. no significant co-channel interference) per metrics available to the RRC 660 from the Channel Equalizer Coefficients Generator 2332 is to assign up to N minus L combinations of an antenna element 652, an Rx-n chain, and a channel MUX receive path to a Complex DFT-n to different frequency channels than the instant link channel to perform DFS or “probe in space” measurements and spectral analysis. For FDD configurations, assigning these combinations to monitor the instant or candidate transmit frequency channels (possibly during a time when the transmitter is otherwise inhibited) can allow the RRC to evaluate potential interference to other conventional PTP links and to adjust transmit resources accordingly. To the extent that the remaining at least L receive chains provide sufficient SNR or SINR to maintain the instant traffic load, this approach allows the RRC 660 to utilize available IBR resources simultaneously for both supporting link traffic and supporting DFS or “probe in space” measurements and spectral analysis.

As described previously, exemplary IBRs advantageously exploit the propagation path diversity usually present in an obstructed LOS environment to send multiple modulated streams concurrently and thus increase overall link throughput. For practical reasons regarding actual field deployments, it is likely that some IBRs will be deployed in locations where propagation may be dominated at least at some times by unobstructed LOS conditions. In such situations, IBR embodiments using the IBR Channel MUX 628 of FIG. 23 may not be able to resolve multiple streams during the training Block 0 of FIG. 26 because all streams arrive at all antennas by substantially similar paths which makes cancellation for multiple demodulation streams impractical. One alternative for the RRC 660 in this situation is to allocate one stream per link only. However, this results in reduced throughput for the link which is not only undesirable but highly counterintuitive for many backhaul operations personnel likely to perceive that unobstructed LOS should be “better” than obstructed LOS. Thus, the RRC 660 may also evaluate certain antenna selection options to intentionally create path diversity in otherwise unobstructed or nearly unobstructed LOS conditions.

A first alternative for the RRC 660 to provide multiple streams with both obstructed and non-obstructed LOS operation is the dynamic testing and possibly selection of mapping different modulator streams to different antenna elements (via separate RF chains) based on different antenna polariza-

tion opposite polarization antenna elements, testing of alternative polarization antenna elements with training data may be pre-arranged in time by RRC control frames exchanged by both IBRs with an instant link. Similarly, the AE-IBR may select any changes in link antenna elements involving polarization and verify an agreed upon changeover time by RRC control frame exchange for reasons analogous to those for RF carrier frequency or channel bandwidth changes. A significant advantage of using polarization diversity amongst the set of selectable antenna elements is that multiple stream throughput can be maintained using a common set of channel equalization techniques as described above for MIMO operation with the exemplary IBR Channel MUX of FIG. 23.

A second alternative for the RRC 660 to provide multiple streams with both obstructed and non-obstructed LSO operation is the dynamic testing and possibly selection of mapping different modulator streams to different antenna elements (via separate RF chains) based on different direction orientations of antenna elements as is possible with exemplary antenna arrays such as those depicted in FIGS. 14 and 15. The RRC 660 in such IBR embodiments can test and/or change-over such multi-directional antenna element combinations entirely at one end of a link without the necessity of exchanging RRC control frames or requiring a coordinated change-over at both IBRs in a link simultaneously. For at least the preceding reason, this second alternative to maintaining multi-stream operation by using multi-directional antenna element combinations that intentionally create propagation path diversity may be desirable for an RE-IBR used in a PMP configuration. An example antenna array suitable for such an RE-IBR is depicted in FIG. 14. Furthermore, this directional orientation diversity strategy to multi-stream operation in otherwise unobstructed LOS is also advantageous compared to the polarization diversity option in PMP deployments because it does not require that the AE-IBR deploy certain resources (e.g., dedicated RF chains and antenna elements) to link modalities of benefit to some RE-IBRs but not others (that do experience obstructed LOS) as may occur with polarization diversity in a PMP deployment.

With reference to FIGS. 7, 32 and 33, the Intelligent Backhaul Management System (IBMS) Agent 700 is an optional element of the IBR that optimizes performance of the instant links at the IBR as well as potentially other IBR links in the nearby geographic proximity including potential future links for IBRs yet to be deployed. As described above in reference to the RRC 660 and depicted in FIG. 32, the primary interaction of the IBMS Agent 700 with the internal elements of the IBR is via a direct bi-directional path to the RRC 660. In one direction, the various policies and configuration parameters used by the RRC 660 to allocate resources within and amongst IBRs with active links to each other are sent from the IBMS Agent 700 to the RRC 660. In the return direction, the RRC 660 reports operational statistics and parameters back to the IBMS Agent 700 both from normal operation modes and from “probe in space” modes as directed by the IBMS Agent 700.

In contrast with the RRC 660, which communicates with other elements of the IBR internally or with other RRC entities at IBRs actively linked to its IBR, the IBMS Agent 700 can receive information from or transmit information to or initiate sessions with other elements of the overall IBMS that are logically located anywhere within any network (subject to appropriate access privileges). As shown in FIG. 32, the IBMS Agent 700 appears to the overall network as an Applications layer entity that exchanges messages typically over TCP/IP and any of the link layer interfaces within the IBR. For the common deployment of IBRs within a cellular system

Radio Access Network (RAN) to backhaul cellular base station sites, it may be necessary for the IBR to tunnel TCP/IP across a cellular RAN specific transport and network layers protocol, such as GPRS Tunneling Protocol (GTP) to reach a gateway that bridges to TCP/IP and on to the desired other IBMS elements.

In some embodiments, the IBMS Agent 700 can act as an autonomous entity that per configuration settings draws information from network resources (whether private or public) and serves as an expert system local to be a specific IBR to optimize its performance. In other embodiments, the IBMS Agent 700 interacts with other peer IBMS Agents at other IBRs within its “interference” zone of influence via self-discovery within the immediate network so that such peer IBMS Agents can collectively optimize IBR performance within the zone. In some embodiments, the IBMS Agent 700 is a client of one or more IBMS servers that may be within the private and/or public networks such that communications with a particular IBMS Agent 700 is always to or from or via such IBMS servers.

The information gathered at the IBR and distilled by the IBMS Agent 700 regarding, for example, operational statistics (such as RSSI, channel equalization metrics, FCS failures, etc.) and resource selections (such as antennas, channel bandwidth, modulator stream assignments, etc.), may be sent to an IBMS server. Such information can be used by the IBMS server to improve performance predictability of future IBR deployments or to enable overall IBR system performance of all links within the purview of the IBMS server by policy optimization across IBRs. The communications from the IBMS server to the IBMS Agents can include such optimized policy parameters based on information from other IBMS Agents and private and/or public databases such as directories of known non-IBR links including their locations, antenna patterns, power levels, carrier frequencies and channel bandwidths, tower heights, etc.

With reference to FIGS. 31 and 32, other exemplary applications layer protocols are shown for the IBR. A Hyper Text Transfer Protocol (HTTP) server application can enable any browser (or browser-like application) on a networked computer, including a WiFi connected mobile device such as laptop, table or smart-phone, to query or modify various IBR configuration parameters. The Simple Network Management System client of FIGS. 31 and 32 is an example of an industry-standard network management utility that can also be used to query or modify IBR configuration parameters using tools such as HP Open View.

The foregoing description of the various elements of the IBR in reference to FIGS. 5-36 have described numerous exemplary embodiments. These exemplary element embodiments may be assembled together in many different combinations and permutations. A very short description of but a few of the possible overall IBR exemplary embodiments is summarized briefly below.

A first exemplary embodiment of an IBR includes the features outlined in

TABLE 1

First Exemplary Intelligent Backhaul Radio Features	
TDD operation	
PTP configuration only (AE-IBR or RE-IBR)	
SC-FDE modulation	
$J_{mod} = J_{dem} = 1$	
$K = 2, M = 4$	
$L = 2, N = 4$	
$Q = 8$, antenna array similar to FIG. 14	

TABLE 1-continued

First Exemplary Intelligent Backhaul Radio Features	
front facet with 2 Vertical and 2 Horizontal polarization antenna elements at 15 dBi	5
side facets each with 1 Vertical and 1 Horizontal polarization antenna elements at 12 dBi	
QPSK, 16 QAM, 64 QAM, 256 QAM	
1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates	
1 or 2 modulated streams	
MMSE combining of up to 4 receive chains	10
MRC weighting of up to 4 transmit chains based on blended weights per chain across entire channel as derived from Rx conditions	
fixed superframe (1 ms) NACK protocol	
up to 28 MHz channel bandwidth	
single MPDU per PPDU	15

This first exemplary embodiment can have a very high MAC efficiency (ratio of MPDU payload bits to overall MAC bits) under heavy loading from the IBR LLC—in excess of 95%. Furthermore, the PHY efficiency (ratio of time where PPDU payload symbols excluding PAD are actually transmitted to superframe time) can exceed 90% for typical channel impairment conditions and ranges of 2 km or less. At 28 MHz symbol rate and channel bandwidth, 256 QAM, 2 modulated streams, 7/8 rate coding, and with MAC and PHY efficiencies of 95% and 90% respectively, the aggregate bi-directional throughput for this first exemplary IBR embodiment can exceed 330 Mb/s with an average end to end latency of about 2 ms.

A second exemplary embodiment of an IBR includes the features outlined in

TABLE 2

Second Exemplary Intelligent Backhaul Radio Features	
TDD operation	35
PMP configuration as the AE-IBR	
OFDM modulation	
$J_{mod} = 2, J_{dem} = 3$	
$K = 4, M = 4$	
$L = 5, N = 5$	40
$Q = 9$, antenna array similar to a 180° azimuth coverage “half-array” of FIG. 15	
per each of 4 facets: 1 Vertical and 1 Horizontal polarization antenna elements at 15 dBi	
on “top”: 1 non-polarized, 180° azimuth (aligned with array) coverage	
8 dBi element connected only to the 5 th Rx chain and 3 rd demodulator core	45
QPSK, 16 QAM, 64 QAM, 256 QAM	
1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates	
up to 2 modulated streams per modulator core	
MMSE combining of up to 5 receive chains	
Eigenbeamforming transmit SDMA of up to 4 transmit chains	50
fixed superframe (0.5 ms or 1 ms) NACK protocol	
TDMA alternate superframe per RE-IBR up to 2 RE-IBRs per modulator core	
up to 28 MHz channel bandwidth	
single MPDU per PPDU	55

This second exemplary embodiment uses OFDM rather than SC-FDE in order to enable transmit SDMA of up to 4 RE-IBRs (only 2 simultaneously) in a highly frequency selective channel. As discussed above, SC-FDE could also be used in this PMP AE-IBR with theoretically similar performance but with more complex baseband processing required. This second exemplary embodiment should have similar MAC efficiency to the first for the 1 ms superframe case but the PHY efficiency (which needs a definition that accounts for data block pilot subchannels and zero-padded subchannels) is typically lower with 85% being excellent. The eigenbeamforming may also require additional overheads in many

53

propagation environments. If 4 RE-IBRs are used in TDMA mode, the latency would expand to 4-5 ms on average for 1 ms superframes and about half that for 0.5 ms superframes. With 2 RE-IBRs that are spatially separable and with operating parameters set as described for the first exemplary embodiment above, aggregate bi-directional throughput for this second exemplary embodiment can be as high as about 600 Mb/s. Note that the “top” antenna which is not connected to a transmit chain can be used to provide the MMSE combiner with an additional degree of freedom to cancel interference when 4 modulated streams from 2 RE-IBRs are being simultaneously received. It can also be used advantageously as a “probe in space” to provide information to the IBMS or to assist in DFS by scanning channels not currently used at the AE-IBR. Also note that although this second exemplary embodiment can be used as an RE-IBR for itself that preferably this role may be filled by the third exemplary embodiment described below.

A third exemplary embodiment of an IBR includes the features outlined in

TABLE 3

Third Exemplary Intelligent Backhaul Radio Features	
TDD operation	
PTP configuration only (AE-IBR or RE-IBR) or PMP configuration as the RE-IBR	
OFDM modulation	
Jmod = Jdem = 1	
K = 2, M = 4	
L = 2, N = 4	
Q = 8, antenna array similar to FIG. 14	
front facet with 2 Vertical and 2 Horizontal polarization antenna elements at 15 dBi	
side facets each with 1 Vertical and 1 Horizontal polarization antenna elements at 12 dBi	
QPSK, 16 QAM, 64 QAM, 256 QAM	
1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates	
1 or 2 modulated streams	
MMSE combining of up to 4 receive chains	
MRC weighting of up to 4 transmit chains based on blended weights per chain across entire channel as derived from Rx conditions	
fixed superframe (1 ms) NACK protocol	
up to 28 MHz channel bandwidth	
single MPDU per PPDU	

Note that while this third embodiment can also be used as a PTP AE-IBR and RE-IBR combination, it is unlikely to provide meaningful performance improvements for PTP compared to the first exemplary embodiment and quite possibly would have slightly lower aggregate bi-directional throughput and would require a less efficient and more expensive power amplifier and stricter phase noise considerations. For commercially-available components today, using OFDM versus SC-FDE at RF carrier frequencies above 10 GHz is extremely challenging. Note that for below 10 GHz operation, it is commercially feasible today to use SDR baseband approaches and common chain and Front-end components to build an IBR software programmable as either the PTP first exemplary embodiment or the PMP RE-IBR third exemplary embodiment.

A fourth exemplary embodiment of an IBR includes the features outlined in

TABLE 4

Fourth Exemplary Intelligent Backhaul Radio Features	
FDD operation	
PMP configuration as the AE-IBR	

54

TABLE 4-continued

Fourth Exemplary Intelligent Backhaul Radio Features	
OFDM modulation	
Jmod = 4, Jdem = 6	
K = 8, M = 8	
L = 10, N = 10	
Q = 18, antenna array similar to FIG. 15	
per each of 8 facets: 1 Vertical and 1 Horizontal polarization antenna elements at 15 dBi	
on “top”: 2 non-polarized, opposite facing, 180° azimuth coverage	
8 dBi elements connected only to the respective ones of the 9 th and 10 th Rx chains and the 5 th and 6 th demodulator cores	
QPSK, 16 QAM, 64 QAM, 256 QAM	
1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates	
up to 2 modulated streams per modulator core	
MMSE combining of up to 10 receive chains	
Closed-loop eigenbeamforming transmit SDMA of up to 8 transmit chains	
fixed superframe (.5 ms or 1 ms) NACK protocol	
TDMA round-robin superframe order per RE-IBR up to 4 RE-IBRs per modulator core	
up to 28 MHz channel bandwidth	
single MPDU per PPDU	

This fourth exemplary embodiment is similar to the second exemplary embodiment except that it utilizes a larger 360° azimuth antenna array and FDD operation as well as 4 time slot per modulator TDMA to support up to 16 RE-IBRs with an aggregate bi-directional throughput of about 2 Gb/s after the increased overhead in efficiencies of the system are accounted for. Latency also increases proportionately if 4 TDMA slots are used.

A fifth exemplary embodiment of an IBR includes the features outlined in

TABLE 5

Fifth Exemplary Intelligent Backhaul Radio Features	
FDD operation	
PMP configuration as the RE-IBR	
OFDM modulation	
Jmod = Jdem = 1	
K = 2, M = 4	
L = 2, N = 4	
Q = 8, antenna array similar to FIG. 14	
front facet with 2 Vertical and 2 Horizontal polarization antenna elements at 15 dBi	
side facets each with 1 Vertical and 1 Horizontal polarization antenna elements at 12 dBi	
QPSK, 16 QAM, 64 QAM, 256 QAM	
1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates	
1 or 2 modulated streams	
MMSE combining of up to 4 receive chains	
Selected unweighted Tx antennas as directed by receiver at opposite end	
fixed superframe (1 ms) NACK protocol	
up to 28 MHz channel bandwidth	
single MPDU per PPDU	

The primary application of this fifth exemplary embodiment is to serve as an RE-IBR for the fourth exemplary embodiment AE-IBR.

A sixth exemplary embodiment of an IBR includes the features outlined in

TABLE 6

Sixth Exemplary Intelligent Backhaul Radio Features	
FDD operation	
PTP configuration only (AE-IBR or RE-IBR)	
SC-FDE modulation	
Jmod = Jdem = 1	
K = M = 2	

TABLE 6-continued

Sixth Exemplary Intelligent Backhaul Radio Features
L = 2, N = 4 Q = 6, antenna array similar to FIG. 13 3 Vertical polarization antenna elements, 2 at 13 dBi and 1 at 18 dBi, 3 Horizontal polarization antenna elements, 2 at 13 dBi, 1 at 18 dBi QPSK, 16 QAM, 67 QAM, 256 QAM, 1024 QAM 1/3, 1/2, 2/3, 3/4, 5/6, 7/8, 9/10 coding rates 1 or 2 modulated streams MMSE combining of up to 4 receive chains selection of 2 antennas in transmit fixed superframe (1 ms) NACK protocol up to 56 MHz channel bandwidth single MPDU per PPDU

This sixth exemplary embodiment provides high aggregate bi-directional throughput of up to about 1.8 Gb/s with moderate complexity relative to other IBRs. This sixth exemplary embodiment performs optimally in propagation channels with only moderate obstructions compared to unobstructed LOS. FDD also provides <1 ms average latency.

A seventh exemplary embodiment of an IBR includes the features outlined in

TABLE 7

Seventh Exemplary Intelligent Backhaul Radio Features
FDD/TDD hybrid operation PTP configuration only (AE-IBR or RE-IBR) SC-FDE modulation Jmod = Jdem = 1 K = 4, M = 8 L = 4, N = 8 Q = 12, antenna array similar to FIG. 14 per each of 3 facets: 2 Vertical and 2 Horizontal antenna elements of 15 dBi each QPSK, 16 QAM, 64 QAM, 256 QAM 1/3, 1/2, 2/3, 3/4, 5/6, 7/8 coding rates up to 4 modulated streams MMSE combining of up to 8 receive chains MRC weighting of up to 8 transmit chains based on blended weights per chain across entire channel as derived from Rx conditions fixed superframe (1 ms) NACK protocol up to 56 MHz channel bandwidth single MPDU per PPDU

This seventh exemplary embodiment has additional resources compared to the sixth exemplary embodiment to provide higher aggregate bi-directional throughput of about 3 Gb/s for a PTP link with <1 ms latency that can operate in a severely obstructed LOS propagation channel. It advantageously uses hybrid FDD/TDD operation wherein each frequency duplexed channel alternates in opposite synchronization to each other between transmit and receive. This enables a relatively straightforward and efficient transmit chain weighting to be derived from receive chain equalization analysis without increasing latency. Furthermore, an additional degree of frequency diversity (and space to the extent different antennas are selected) is achieved. The FDD/TDD hybrid can be utilized on any FDD IBR deployment where spectrum regulations permit it. To the extent each FDD band operation relies on band-specific band-select filters in the Front-ends, then additional circuit complexity for switching between transmit and receive is needed.

Note that the preceding embodiments are a small subset of the possible IBR embodiments that are enabled by the disclosure herein. Note further that many additional optional structures or methods described herein can be substituted within the above exemplary embodiments or other embodiments.

For example, TDD CSMA could be used advantageously in high interference spectrum allocations or where required by spectrum regulations as a substitute for fixed superframe timing in the above exemplary embodiments.

Note also that all of the above exemplary embodiments are compatible with any RF carrier frequencies in the range of interest from approximately 500 MHz to 100 GHz.

Note further that in multi-channel embodiments, it is possible to use different access and MAC protocols in different channels especially where advantageous for or required by spectrum regulations. For example, an IBR link may advantageously provide a “base” throughput capability in a channel expected to have minimal interference but for which licensing costs or regulatory restrictions limit total throughput. Then a “surge” throughput capability can be provided in a second channel, such as unlicensed spectrum, where throughput can be higher but the risk of temporal interference outages is also higher.

As evident from the above exemplary embodiments, OFDM is typically used for PMP deployments because at a baseband processing level it is relatively less complex than SC-FDE if frequency-selective channels are to be used with transmit SDMA (at least at the AE-IBR). However, OFDM has higher peak to average ratio and is more sensitive to carrier frequency offset and phase noise than SC-FDE. Thus, for PTP, SC-FDE is often preferable, especially for operation at RF carrier frequencies above 10 GHz where commercially viable components are expensive for either OFDM power amplification or OFDM-compatible local oscillator specifications.

Note that in backhaul applications where links are nominally continuous, additional techniques can improve PHY efficiency. For example, with OFDM, the training block can be combined with a PLCP Header block by interleaving sub-channels appropriately. This is also possible for SC-FDE if DFT pre-coding (i.e. Tx block Assembler-k includes a DFT and Tx-Mux-m includes an IDFT) is used. DFT pre-coding for SC-FDE can also be used to pulse shape the transmitted waveform similar to OFDM zero padding and/or windowing. The training block in SC-FDE can also be shorter than the data blocks to save PHY overhead by either using another FFT for training or switching the FFT bin size during the training block. The channel equalization function so derived is then interpolated for use on the longer data blocks with additional frequency bins. Also, when using DFT pre-coding in SC-FDE, it is possible to time multiplex an FFT block between transmit and receive or within transmit or receive such that all four FFT operations of an SC-FDE transceiver can be realized by 2 or even 1 FFT hardware core. Another technique to simplify PMP deployment of IBRs is to use OFDM in the forward link from AE-IBR to the multitude of RE-IBRs, and then use SC-FDE in the reverse link from RE-IBR to AE-IBR. This enables the advantages of transmit eigenbeamforming at the AE-IBR in a frequency-selective channel based primarily on receive equalization while keeping the RE-IBRs relatively straightforward with much simpler transmitters than in the OFDM only case.

Numerous additional variations of the above-described elements of the IBR can also be advantageously utilized in substitution for or in combination with the exemplary embodiments described above. For example, antenna elements need not always be shared between transmit and receive whether in TDD or FDD mode. In certain embodiments, it is preferable to have a smaller number of transmit antenna elements, often with broader azimuthal coverage than those used by the receive antenna elements, that are always used versus a selectable larger number of receive

antenna elements. In some embodiments with separate transmit and receive antenna elements, the respective front-ends of FIGS. 11 and 12 would not require respectively the SPDT switch or the band selection within a duplexer filter and either the receive path elements for a front-end coupled to a transmit antenna element or the transmit path elements for a front-end coupled to a receive antenna element, which further advantageously reduces cost and complexity of the IBR. In such embodiments, each transmit antenna element is coupled to a respective transmit power amplifier which is in turn coupled to a respective transmit RF chain wherein such couplings may be selectable RF connections or fixed RF connections—the latter for IBR embodiments wherein certain transmit RF chains are always connected to certain transmit power amplifiers and certain transmit antenna elements whenever the IBR is in a transmit mode. Similarly, each receive antenna element is coupled to a respective receive low noise amplifier which is in turn coupled to a respective receive RF chain wherein such couplings are typically selectable RF connections as described herein.

As another example, the NACK protocol described above with reference to FIG. 33 and various IBR embodiments can be extended to either ACK/NACK, single NACK as described above, or persistent NACK until delivered. Furthermore, such exemplary choices can be applied to individual MSDUs based on either an indicator from the IBR LLC of FIG. 33 or inspection of MSDU class of service or type of service header field bits as defined by a policy either within the IBR MAC or updateable via the optional IBMS Agent 700. For example, MSDUs corresponding to Voice over Internet Protocol (VoIP) packets are typically sent at a high class of service priority but with so little tolerance for latency that it may be preferable to send such MSDUs with no ACK/NACK retransmission option or with only the single NACK protocol described previously. At the opposite extreme, certain data transfers, such as for example only, cellular network control or management plane messages or user file transfer data, may tolerate considerable or unpredictable latencies associated with persistent retransmission until NACK=0 rather than rely on much slower upper layer retransmission protocols. The policy for mapping MSDUs to a particular ACK/NACK strategy may also be responsive to radio channel conditions and/or loading as determined within the various elements of the IBR described above. For example, when the current packet failure rate is very low and/or the loading demand on the IBR is high compared to the capacity of the MCS, then one policy may be to minimize retransmissions by transmitting most MSDUs with no ACK/NACK or a single NACK. Alternatively, for a high packet failure situation and/or low demand at a given MCS, the opposite strategy may be used. Also, any of the IBR embodiments described herein that use copper Ethernet interfaces may also use such interfaces to supply Power over Ethernet (PoE) to the IBR.

One exemplary embodiment of an antenna system for the IBR is based on a combination of angle-beam diversity and polarization diversity. In some embodiments, an advantageous far-field radiation pattern may be achieved with multiple instances of a two-port, dual-polarity sector antenna arranged side-by-side at various angles.

An exemplary antenna is illustrated in FIG. 37. The corresponding overlaid far-field radiation patterns for the antenna of FIG. 37 are illustrated in FIG. 38. It should be noted that the far-field patterns are identical for both polarizations, and for convenience are shown utilizing a single plot for both polarizations. It should be noted that such convention is utilized where applicable in subsequent diagrams and plots as well.

In FIG. 37, each antenna is used as bi-directional (i.e., for both transmit and receive). Each antenna may be used in a time-division-duplex system or frequency division duplexed system. Signals 3710 (A,B,C,D) are respectively transmitted and received by two port, dual polarized sector antennas 3701 (A,B,C,D), which are coupled to base band signal processing 3730 respectively by transceivers 3720 (A,B,C,D). It will be appreciated that the antenna system of FIG. 37 may be employed any and all techniques disclosed herein.

FIG. 38 is a plot of an exemplary far field antenna pattern of a sector antenna panel. Patterns 3801(A,B,C,D) respectively correspond to sector antennas 3701(A,B,C,D). The bore sight of the antenna system (corresponding to zero degrees alignment) of FIG. 37 corresponds to bore sight 3820 in FIG. 38, with each incremental radial reference corresponding to 30 degree increments. In FIG. 38, increasing reference gain increments 3815, 3810, 3805 correspond to steps of 5 dB.

FIG. 39 illustrates an alternative sector antenna panel and processing, including separate transmit and receive antennas. The arrangement of FIG. 39 includes dedicated unidirectional transmit/receive antennas, which allows the patterns to be set independently. Dedicated unidirectional transmit/receive antennas can be advantageous in certain types of regulated bands, such as EIRP limited bands. The arrangement shown in FIG. 39 is similar to the arrangement of FIG. 37, but an additional two-port, dual polarity transmit antenna 3950 is provided as a forward-facing center panel. In FIG. 39, the four other two-port dual polarity antennas 3901 (A,B,C,D) are used as receive-only antennas.

FIG. 40 is a plot of an exemplary far field antenna pattern of the alternative sector antenna panel including separate transmit and receive antennas of FIG. 39. The transmit sector panel 3950 has a wider azimuthal beam width so that the transmit beam width 4030 is substantially equal to the composite receive antenna beam width 4001(A,B,C,D).

FIG. 41 illustrates an alternative sector antenna panel and processing including separate transmit and receive antennas. The antenna panel of FIG. 41 has an alternative geometric arrangement of elements. In FIG. 41, the order of the receive antennas 3901(A,B,C,D) in the antenna array is re-assigned. In FIG. 39, the inner-most receive antennas (3901B and C) are arranged to the outside of the array (4101A,D) and increased in forward angular orientation such that the overlap of the antenna patterns relating to antennas (4101A,D) is increased. In such an arrangement, the antennas 4101A,B,C,D correspond respectively to the antenna patterns of FIG. 42 is a differing order of antenna 4201B,A,D,C due to the aforementioned modifications in arrangement and orientation. Spatial diversity is achieved due to the additional physical distance between antennas 4101A and 4101D despite the overlap of the antenna patterns 4201 B and 4201C. The transmit antenna pattern of antenna 4150 is identical in FIG. 42 to the pattern of antenna 3950 as shown in plot 4030.

FIG. 43 illustrates a dual-polarity, two-port patch antenna element including feed and grounding points. The arrangement of FIG. 43 may be used as a component of the antenna arrays described herein for use as a receive or a transmit antenna panel. In FIG. 43, the dual-port design is based on a circular patch antenna 4300. In FIG. 43, two orthogonal modes are excited by two orthogonal probe feeds 4320 and 4310. Each mode excites linearly polarized far-field radiation. A shorting pin 4330 is provided in the center of the patch to suppress the DC-mode of the patch that would normally be the primary mechanism creating undesired coupling between the two ports. The construction is based on two microwave-grade PCBs: one is used as a ground-plane 4302 for the patch,

59

and the other contains the etched circular patch **4301**. The ground-plane PCB **4302** also provides micro strip feed structure to feeds **4320** and **4310**.

FIG. **44A** and FIG. **44B** illustrate an exemplary dual-polarity, two port, patch antenna array. The antenna array includes etched patch antenna elements **4401A-H**. In FIGS. **44A** and **44B**, the antenna elements **4401A-H** include patch **4301**. A shared patch array ground plane **4402**, which corresponds to the ground plane **4302**, is provided for each patch element. Each patch element **4401A-H** includes a two port antenna element utilizing orthogonal polarization modes. In one embodiment, feeds **4410A-H** are provided for a first polarization, and feeds **4420A-H** are provided for the second polarization.

In some embodiments, the collective patch antenna array **4400A** provides for a collective two port interface and provides for a common micro strip cooperate feed network integrated with the ground plane **4403** PCB. For example, a first port feeds each of the same polarization feeds **4410A-H** providing for a polarized sub-array of array **4400A**, while a second port feeds each of the other same polarization feeds **4420A-H** providing for the other polarized sub-array. In some embodiments, a micro strip cooperative feed network is provided for each polarization to have a common delay from each array port to each of the respective polarization feeds to achieve a desired array factor, and, in turn, a desired array far-field pattern. In some embodiments, variations on the relative array port to polarization feeds are provided to achieve various modified antenna patterns such as a modified beam width, side lobe levels, or the like.

In some embodiments, the antenna array also includes a grounded fence structure **4440A** and **4440B** which provides advantages relating to improved azimuthal array gain directivity, and side lobe levels, as well as potentially improved near field isolation from other additional antenna structures. In some embodiments, the antenna array includes supporting structures **4405A-J** to provide structural support of printed patch PCB **4402** from ground plane **4403** PCB. It will be appreciated that the patch array **4400A** may be used in the receive or transmit panels of FIG. **41**, as well as other embodiments disclosed herein.

FIG. **45A** and FIG. **45B** illustrate an exemplary single-polarity, single port, printed dipole antenna element **4500**. The antenna element **4500** may be used in an antenna array. FIG. **45A** and FIG. **45B** show more than a single representative element; however, only a single element **4500** is described with reference to FIGS. **45A** and **45B**.

In FIGS. **45A** and **45B**, the dipole antenna element **4500** is etched on printed circuit board **4502**. A first arm of the element **4510** is connected to a micro strip feed network integrated with PBC **4502** via a BALUN structure element **4514**. The BALUN structure includes structures **4514** and **4507**. Structures **4514** and **4507** receive a unbalanced signal from the micro strip feed network **4516** utilizing post **4515**. The associated post **4515** couples a signal to the BALUN element **4514**, which produces a balanced signal (together with BALUN element **4507**), which then is coupled to balanced impedance matching structures and **4511** and **4506**. The post passes through ground plane **4501** via a clearance structure **4503**. An element arm **4505** acts as a counter poise to element arm **4510**, and connects to BALUN structure **4507** via impedance matching **4506**. A ground plane **4501** is provided to provide additional directive gain in the Z axis. Collectively, the printed dipole element **4500** provides a single port, single linear polarization directive antenna radiation pattern. In some embodiments, the BALUN ground plane **4507** is electrically connected to the reflector ground plane **4501** to ensure

60

a low-impedance transition from the backplane microstrip structure to the antenna microstrip structure. In the embodiment shown in FIGS. **45A** and **45B**, the arms of the printed dipole are swept back primarily to widen the beamwidth of each element. A secondary benefit of swept back arms is that the impedance bandwidth (VSWR bandwidth) of the antenna is improved.

FIG. **46A** and FIG. **46B** illustrate an exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements. In some embodiments, the antenna elements **4605A-J** are equivalent to the exemplary single-polarity, single port, printed dipole antenna element **4500** shown in FIG. **45A** and FIG. **45B**. Elements **4605A,C,E,G**, and **I** are rotated at a positive 45 degrees, providing for a positive "slant 45" polarization "sub-array." The elements **4605A,C,E,G**, and **I** are commonly fed with a shared cooperate feed network and are coupled to a first port of array **4600A**. Elements **4605B,D,F,H**, and **J** are rotated at a negative 45 degrees, providing a negative "slant 45" polarization sub-array. The elements **4605B,D,F,H**, and **J** provide for the opposite orthogonally polarized sub-array. Elements **4605B,D,F,H**, and **J** are commonly fed with another shared cooperate feed network and are coupled to the other second port of array **4600A**. In some embodiments, the positive 45 degree slant dipoles are placed offset from the negative 45 degree slant dipoles to minimize the mutual coupling from any of elements **4605B,D,F,H**, and **J** to any nearby elements in the set **4605A,C,E,G**, and **I**. Additionally, the placement of the elements **4605B,D,F,H**, and **J** are the nulls of elements **4605A,C,E,G**, and **I** and vice versa. Such an arrangement is achieved by placing the center of each element of one column aligned at 90 degrees to dipole arms of the opposite column, and having appropriate offset spacing so as to not couple the elements.

FIG. **46C** illustrates the printed circuit board providing for a ground plane of an exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements, showing a cooperate feed network for each single polarized sub-array associated with each port. The antenna array includes a first array port **4640** to couple a signal for a first polarized sub-array utilizing micro-strip lines and divider structures. The antenna array includes a second port **4620** to couple a signal for a second orthogonal polarization utilizing micro-strip lines and divider structures.

In some embodiments, the group delay from port **4640** to each associated antenna element feed posts **4615B,D,F,H,J** is equalized. In some embodiments, each feed network has equalized electrical lengths such that each port is excited in-phase. Likewise, in some embodiments, the group delay from port **4620** to each associated antenna element feed posts **4615A,C,E,G,I** is designed to also be equalized. In some embodiments, the signal power coupled to or from each of the array ports **4640**, **4620** with each of the associated antenna sub-array elements is equal. Such splitting or combining and delay equalization is performed utilizing micro-strip dividers and impedance matching structures, such as **4660**, **4655**, **4652**, **4645**, **4650**, **4625**, and **4630**. In some embodiments, the delays of each of the feed networks are equal to each other, or have a pre-determined known different. Such embodiments provide for a reduced complexity when performing specific transmission or reception beam forming operations, or adaptive polarization operations. Additionally, such offsets may be compensated for in digital baseband, or in other RF delay equalization components. One such structure for adding delay or phase offsets is the transmit channel equalizer Tx-CE-m **2308**. Additionally, in some embodiments, each polarizer sub-array (fed by ports **4640** and **4620** respectively) is fed

61

with separate transmitter chains utilizing different transmit channel equalizers **2308** within the IBR Channel MUX **628** as shown in FIG. **23A**.

In some embodiments, adjustable polarization is provided by defining the relationship of the respective Tx-CE **2308**

weights $\overrightarrow{WT}_{m,k}$ feeding each of the polarized sub-arrays of the same array panel as applied to the transmit symbol stream

$\overrightarrow{TxBlock_k}$. For example, transmitting a Tx Block solely from one polarized sub-array relative to the other results in the respective polarization alone. However, transmitting a Tx Block with the same phase from each of the cross polarized sub-arrays results in a polarization which is linear and at 45 degrees between the two respective polarizations. Further adjustment of the “slant” of the resulting polarization may be achieved utilizing various relative amplitudes between the two equalized transmit blocks.

Additional TX polarizations may be achieved by adjusting the relative amplitude and phase between the same transmit

block $\overrightarrow{TxBlock_k}$ on each of the cross polarized sub-arrays of the same panel. As an additional example, circular polarized transmission may be achieved by delaying the phase of one transmitted signal by 90 degrees for right handed polarization (or -90 degrees for left handed circular). Such adjustments may be made adaptively based upon various criteria, including metrics feedback from the target receiving IBR. Such processing may be extended to antenna arrays and sub-arrays which are not within the same array or panel. Such processing results in transmitter beam forming operations. Further, each

transmit streams or transmit block $\overrightarrow{TxBlock_k}$ may be adjusted independently as described in relation to embodiments of FIG. **24** and FIG. **24A**, and may provide for frequency selective block specific adaptive or adjustable polarization transmission or other forms of beam forming.

FIG. **47A** illustrates a plot of an exemplary far field elevation antenna pattern of a dual-polarity, two port, antenna array utilizing printed dipole antenna elements, and FIG. **47B** illustrates a plot of an exemplary far field azimuthal antenna pattern of a dual-polarity, two port, antenna array utilizing printed dipole antenna elements. In some embodiments, the antenna patterns for each of the polarized sub-arrays provide substantially the same pattern; accordingly, only one representative plot is shown representing both sub-arrays.

FIG. **48** illustrates an alternative exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements, including alternative printed dipole antenna elements and further providing vertical and horizontal polarizations. Array **4800** provides for vertically and horizontally polarization, rather than the slant polarizations of **4600A**. The arrangement of the elements in the alternative array configurations of **4800** are modified to minimize the mutual coupling from any of the elements **4810A,B,C,D**, and E to any of the elements are within the antenna pattern nulls of elements **4820A,B,C**, and D. There is strong mutual coupling between the elements in each subset (**4810A,B,C,D**, and E) and (**4820A,B,C**, and D), but this coupling is not detrimental to the operation of a uniformly excited (both amplitude and phase) antenna array, but the opposite is not true in this arrangement.

Elements **4820A,B,C,D** are within the antenna pattern of the elements of the orthogonally polarized sub array **4810A,B,C,D** and E. However, because elements **4810A,B,C,D** and E are oriented at 90 degrees to the field lines of the elements **4820A,B,C,D**, they do not substantially interact electromagnetically. The parasitic elements **4905A** and B increase the azimuthal beamwidth of the horizontal polarized antenna

62

element beyond that achievable with a conventional dipole. The parasitic elements **4905A** and B also improve the impedance bandwidth (VSWR bandwidth) of the antenna element, and in turn, the bandwidth of the array.

FIG. **49A** illustrates an alternative exemplary printed dipole antenna element. In some embodiments, antenna **4900A** includes elements **4810A,B,C,D** and E. Antenna **4900A** includes a printed circuit board **4901** that is etched to provide two branches of a dipole antenna arms **4902** and **4903**, balanced feed ports **4925A**, **4925B**, directors **4905A**, **4905B** and securing tabs **4915A** and **4915B**.

In some embodiments, the length and position of the parasitic directors **4905A**, **4905B** relative to the driven dipole is optimized to achieve a desired azimuthal radiation bandwidth. The position may be optimized by varying the distance between the end of the dipole branches **4903**, **4902** and the parasitic elements **4905A**, **4905B**, as well as the fore and aft offset of the parasitic element **4905A**, **4905B** relative to the dipole branches **4903**, **4902**.

The antenna also includes tabs **4910A,B** and **4920A** and B, which are used to secure the antenna element to the ground plane PCB **4801**, and are soldered in place as shown in FIG. **49B**. Each tab mates with a plated through hole slot in the PCB **4945**, which allows for a securing solder joint. In some embodiments, the soldering is performed with a wave soldering operation. In some embodiments, the soldering is performed using a solder paste reflowed, as typically used in surface mount operations, and wicked into the associated joints. In some embodiments, adhesive or epoxy is used to form one or more of the tab PCB joints.

In some embodiments, a slot **4935** is provided to provide a mechanical connection with an alternative implementation of antenna elements **4820A,B,C**, and D, as shown in FIG. **50** and FIG. **51**. Antenna elements **4820A,B,C** and D may be fabricated on a single printed circuit board **5001**. Slots **5005A-E** are provided to mate with slot **4935** of antenna **4900A** for mechanical rigidity. Tabs **5015A-D** and **5025A-D** are secured to PCB **5110**, as described with respect to FIGS. **49A** and B. Printed dipole element arms **5010A-D**, and **5020A-D** are etched in printed circuit board **5001** to form the antenna sub-array structure **5000**.

FIG. **51** illustrates an exemplary assembly a dual-polarized, two port, printed dipole antenna array utilizing a printed dipole structure. Printed circuit board **5101** receives tabs associated with sub-array board **5000** and elements **4900A** as described above. In some embodiments, the vertical element spacing is 0.65 of a free-space wavelength (equivalent to 35 mm). In FIG. **51**, the vertical polarized elements are not swept back to minimize elevation beam width. The vertical elements feature flared out width on the dipole branches to improve the impedance bandwidth (allowing for a similar impedance bandwidth to that of the horizontal elements). In the elevation plane, the beam width of the vertical element pattern is narrower than the beam width of the horizontal element pattern. This difference in element pattern beam width is compensated in the array factors. In some embodiments, the vertical polarization array uses four elements and the horizontal polarization array uses five elements. The extra element in the horizontal array helps to increase the horizontal array factor (i.e. narrow the pattern beam width).

FIG. **52A** illustrates an exemplary horizontally arranged intelligent backhaul radio antenna array, FIG. **52B** illustrates an alternative view of an exemplary horizontally arranged intelligent backhaul radio antenna array, and FIG. **52C** illustrates a top of an exemplary horizontally arranged intelligent

63

backhaul radio antenna array. Embodiments the arrangement of FIGS. 52A-C include the features of the antenna array of FIG. 41.

In some embodiments, the receive arrays 5210A-D are configured in an arrangement of FIG. 41, wherein the dedicated two port dual polarized receive arrays 4101A-D are implemented as the exemplary dual-polarity, two port, patch antenna array 4400A as shown in FIG. 44A and FIG. 44B. The dedicated transmit array of FIG. 52A may be implemented as the exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements 4600A of FIGS. 46A,B, and C, with the addition of tabs 5201 and slots 5205. Such tabs and slots provide for the shaping of the directive patterns of one column of polarized elements (e.g., sub-array elements 4605A,C,E,G,I) relative to the other column (e.g., sub-array elements 4605B,D,F,H,J). The tabs on each side of the array 4600A allow for the patterns of the two offset transmitter antenna columns to have substantially centered and equal radiation patterns. It will be appreciated that if no slots are utilized, the ground plane would appear larger on one side of the columns than on the other side of a specific column.

Exciting the dipoles 4605 with an applied RF signal causes induced surface currents on the backplane 4600A which are oriented in the same direction as the source dipole. The slots closest to each column are parallel with the surface currents and do not significantly impede the current flow. The slots located at a further distance from the dipoles are oriented orthogonally to the surface currents, causing the ground plane to look effectively narrower on the far side. The slots and tabs allow for the oppositely aligned antenna elements to not be effected by the "wider" ground plane of the tabs, where the "aligned" tabs will appear as ground to the aligned (the nearer) element column. Such an arrangement substantially equalizes the radiation patterns of each of the polarized sub-arrays. For example, sub array antenna elements 4605B,D,F,H, and J of 4600A, as shown in FIG. 46A, are rotated at a negative 45 degrees, providing for a negative "slant 45" polarization sub-array. The elements 4605B,D,F,H, and J are aligned with tab 5201 and slot 5205. The area of the tabs on the side of the ground plane is advantageous in effecting the radiation pattern because the features appear as a solid ground to the sub-array elements 4605B,D,F,H, and J and appear open-circuited to the oppositely aligned sub-array elements 4605A,C,E,G,I.

The receive two port, dual polarized patch antennas arrays 5210A,B,C, and D correspond to the idealized antenna patterns of FIG. 42, respectively, as 4201B,A,D,C. The transmit pattern of the two port, dual polarization printed dipole array of 4600A was shown in relation to FIG. 47A and FIG. 47B, and is further shown with an associated coverage as 4030 in FIG. 40, with respect to such a transmit pattern substantially matching the cumulative receive patterns of arrays 5210A,B, C, and D. Additionally, the elevation patterns of 5210A,B,C, and D substantially match the elevation patterns of 4600A. As noted, each array 5210A,B,C,D and 4600A includes two ports coupling to associated oppositely polarized sub-arrays. The antenna array includes a radome 5220, which may be made of any radiolucent material as known in the art.

FIGS. 53A,B, C and D illustrate an exemplary vertically arranged intelligent backhaul radio antenna array. The antenna array of FIGS. 53A-D physically shifts the receiver arrays from FIG. 52A-C so that, rather than being linearly aligned horizontally, two of the receive arrays are stacked above the transmitter array and two are stacked below the transmitter array in a vertical configuration. Variations of this

64

embodiment may include all four receive arrays being in a vertical alignment, or other combinations.

The two port dual polarized patch array 4400A is utilized with dedicated receive arrays 5330A,B,C, and D. In contrast to the antenna array of FIGS. 52A-C, the transmitter array 4600A is replaced in FIGS. 53A-D with the alternative exemplary dual-polarity, two port, antenna array utilizing printed dipole antenna elements, including alternative printed dipole antenna elements and further providing vertical and horizontal polarizations 4800. Either transmit array or other variations may be utilized with any or all of the exemplary antenna array embodiments disclosed herein.

In some embodiments, the total number of receiver sub-arrays which are available for use in the reception of signal, may exceed the number of IBR receivers present in a specific radio configuration. As discussed in relation to FIG. 10 an IBR RF Switch Fabric 1012 may be utilized couple selectable sub-arrays to specific receivers RF-RX-1 to RF-RX-N. It will be appreciated that any and all variations disclosed herein may be used with the embodiments of FIGS. 52A-C. In addition, the embodiments of FIGS. 13-15 may be used with the embodiments of FIGS. 52A-C and the embodiments of FIGS. 53A-D.

As explained above, in some embodiments, the Switch Fabric 1012 of IBR RF 628 may be controlled by the radio resource controller (RRC). It will be appreciated that the functions and capabilities of the RRC disclosed herein may be used with the embodiments of FIGS. 52A-C and FIGS. 53A-D.

In some embodiments, the delays of each of the two feed networks are equal to each other, or are a pre-determined known different. Such embodiments provide for a reduced complexity when performing specific transmission or reception beam forming operations, or adaptive polarization operations. Additionally such offsets may be compensated for in digital baseband, or in other RF delay equalization components. One such structure for adding delay or phase offsets is the transmit channel equalizer Tx-CE-m 2308. Additionally, in some embodiments, each polarizer sub-array (fed by ports 4640 and 4620 respectively) is fed with separate transmitter chains utilizing different transmit channel equalizers 2308 within the IBR Channel MUX 628 as shown in FIG. 23A.

Adjustable polarization may be provided by defining the relationship of the respective Tx-CE (2308) weights

$\overline{WT}_{m,k}$ feeding each of the polarized sub-arrays of the same array panel as applied to the transmit symbol stream

$TxBlock_k$. For example, transmitting a Tx Block solely from one polarized sub-array relative to the other results in the respective polarization alone. However, transmitting a Tx Block with the same phase from each of the cross polarized sub-arrays results in a polarization which is linear and at 45 degrees between the two respective polarizations. Further adjustment of the "slant" of the resulting polarization may be achieved utilizing various relative amplitudes between the two equalized transmit blocks.

Additional TX polarizations may be achieved by adjusting the relative amplitude and phase between the same transmit

block $TxBlock_k$ on each of the cross polarized sub-arrays of the same panel. As an additional example, circular polarized transmission may be achieved by delaying the phase of one transmitted signal by 90 degrees for right handed polarization (or -90 degrees for left handed circular). Such adjustments may be made adaptively based upon various criteria, includ-

65

ing metrics feedback from the target receiving IBR. For example, the target IBR may monitor any number of receiver metrics, including, for example, receive signal RSSI, C/I, E_B/N_o , E_B/I_o , FER, BER, propagation channel state information (CSI), MIMO stream to stream channel conditioning (ease of spatial multiplexing and channel correlation matrix conditioning), receive polarization or angular spectrum relative to interferer polarization or angular spectrum, or the like. All, some, or a derived metric from a target IBR may be feedback to the transmitting IBR, and/or derived directly for TDD or other reciprocal channel duplexing approaches. Such metrics may be used to calculate, derive, or adaptively determine the transmission equalizer weights of 2332 for frequency flat or frequency selective transmission equalization approaches. Such processing may be extended to antenna arrays and sub-arrays which are not within the same array or panel. Such processing may result in transmitter beam forming operations. Further, each transmit streams or transmit block $TxBLOCK_k$ may be adjusted independently as described in relation to embodiments of FIG. 24 and FIG. 24A, and may provide for frequency selective block specific adaptive or adjustable polarization transmission or other forms of beam forming.

One or more of the methodologies or functions described herein may be embodied in a computer-readable medium on which is stored one or more sets of instructions (e.g., software). The software may reside, completely or at least partially, within memory and/or within a processor during execution thereof. The software may further be transmitted or received over a network.

The term "computer-readable medium" should be taken to include a single medium or multiple media that store the one or more sets of instructions. The term "computer-readable medium" shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by a machine and that cause a machine to perform any one or more of the methodologies of the present invention. The term "computer-readable medium" shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media.

Embodiments of the invention have been described through functional modules at times, which are defined by executable instructions recorded on computer readable media which cause a computer, microprocessors or chipsets to perform method steps when executed. The modules have been segregated by function for the sake of clarity. However, it should be understood that the modules need not correspond to discreet blocks of code and the described functions can be carried out by the execution of various code portions stored on various media and executed at various times.

It should be understood that processes and techniques described herein are not inherently related to any particular apparatus and may be implemented by any suitable combination of components. Further, various types of general purpose devices may be used in accordance with the teachings described herein. It may also prove advantageous to construct specialized apparatus to perform the method steps described herein. The invention has been described in relation to particular examples, which are intended in all respects to be illustrative rather than restrictive. Those skilled in the art will appreciate that many different combinations of hardware, software, and firmware will be suitable for practicing the present invention. Various aspects and/or components of the described embodiments may be used singly or in any combination. It is intended that the specification and examples be

66

considered as exemplary only, with a true scope and spirit of the invention being indicated by the claims.

What is claimed is:

1. A backhaul radio comprising:

a horizontally arranged backhaul radio antenna array comprising:

a plurality of directional receive antenna arrays, wherein each of said plurality of directional receive antenna arrays comprises a plurality of receive antenna sub-arrays, and wherein each of said plurality of receive antenna sub-arrays comprises a plurality of receive antenna elements; and

a directional transmit antenna;

one or more demodulator cores, wherein each demodulator core is capable of demodulating one or more of a plurality of receive symbol streams to produce one or more receive data interface streams;

a plurality of receive radio frequency (RF) chains, wherein each receive RF chain is capable of converting from one of a plurality of receive RF signals to a respective one of a plurality of receive chain output signals; and

a frequency selective receive path channel multiplexer, interposed between the one or more demodulator cores and at least two of the plurality of receive RF chains, wherein the frequency selective receive path channel multiplexer is capable of generating the plurality of receive symbol streams from at least two of the plurality of receive chain output signals provided by said at least two of the plurality of receive RF chains;

wherein each of the plurality of receive antenna sub-arrays is coupled or couplable to at least one receive RF chain; and

wherein all of the plurality of directional receive antenna arrays are horizontally offset relative to the directional transmit antenna.

2. The backhaul radio of claim 1, wherein the frequency selective receive path channel multiplexer comprises at least one of a Space Division Multiple Access (SDMA) combiner or equalizer, a maximal ratio combining (MRC) combiner or equalizer, a minimum mean squared error (MMSE) combiner or equalizer, an Eigen Beam Forming (EBF) combiner or equalizer, a receive beam forming (BF) combiner or equalizer, a Zero Forcing (ZF) combiner or equalizer, a channel estimator, a Maximal Likelihood (DL) detector, an Interference Canceller (IC), a VBLAST combiner or equalizer, a Discrete Fourier Transformer (DFT), a Fast Fourier Transformer (FFT), or an Inverse Fast Fourier Transformer (IFFT).

3. The backhaul radio of claim 1, wherein each one of the one or more demodulator cores comprises at least a decoder and a soft decision symbol demapper.

4. The backhaul radio of claim 1, wherein each one of the plurality of receive RF chains comprises at least a vector demodulator and two analog to digital converters that are capable of producing the respective one of the plurality of receive chain output signals, each said respective one of the plurality of receive chain output signals comprised of digital baseband quadrature signals.

5. The backhaul radio of claim 1, further comprising:

at least one additional directional transmit antenna, wherein each one of the directional transmit antenna and the at least one additional directional transmit antenna is a directional transmit antenna array, wherein each of said directional transmit antenna arrays comprises a plurality of transmit antenna sub-arrays, and wherein each of said plurality of transmit antenna sub-arrays comprises a plurality of transmit antenna elements;

67

one or more modulator cores, wherein each one of the one or more modulator cores is capable of modulating one or more transmit data interface streams to produce one or more of a plurality of transmit symbol streams; and a plurality of transmit RF chains, wherein each one of the plurality of transmit RF chains is capable of converting from one of a plurality of transmit chain input signals to a respective one of a plurality of transmit RF signals each comprising at least one of said plurality of transmit symbol streams;

wherein each of said plurality of transmit antenna sub-arrays is coupled or couplable to at least one transmit RF chain; and

wherein all of the plurality of directional receive antenna arrays are horizontally offset relative to all of said directional transmit antenna arrays.

6. The backhaul radio of claim 5, wherein the backhaul radio is capable of radiating at least one transmit RF signal from at least one of the plurality of transmit antenna sub-arrays at the same time that said at least two of the plurality of receive RF chains are capable of providing said at least two of the plurality of respective receive chain output signals to the frequency selective receive path channel multiplexer.

7. The backhaul radio of claim 5, wherein an antenna pattern beamwidth of each of the receive antenna sub-arrays is similar in elevation to an antenna pattern beamwidth of each of the transmit antenna sub-arrays in elevation.

8. The backhaul radio of claim 7, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays is narrower in azimuth than an antenna pattern beamwidth of each of the transmit antenna sub-arrays in azimuth.

9. The backhaul radio of claim 7, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays is narrower in elevation than an antenna pattern beamwidth of each of the receive antenna sub-arrays in azimuth.

10. The backhaul radio of claim 5, wherein each of the plurality of transmit antenna elements is one of a patch antenna element or a printed dipole antenna element.

11. The backhaul radio of claim 5, wherein each one of the one or more modulator cores comprises at least an encoder and a symbol mapper.

12. The backhaul radio of claim 5, wherein each one of the plurality of transmit RF chains comprises at least two digital to analog converters and a vector modulator that are capable of operating on the respective one of the plurality of transmit chain input signals, each said respective one of the plurality of transmit chain input signals comprised of digital baseband quadrature signals.

13. The backhaul radio of claim 5 further comprising:

a transmit path channel multiplexer, interposed between the one or more modulator cores and the plurality of transmit RF chains, to produce the plurality of transmit chain input signals provided to the plurality of transmit RF chains from the plurality of transmit symbol streams, wherein the transmit path channel multiplexer operates on one or more successive blocks of symbols from the plurality of transmit symbol streams.

14. The backhaul radio of claim 13, wherein the transmit path channel multiplexer comprises:

at least one cyclic prefix adder that prepends or appends a fraction of a block of transmit chain time domain samples to be transmitted sequentially in time to such block of transmit chain time domain samples to produce a cyclically-extended block of transmit chain time domain samples.

68

15. The backhaul radio of claim 13, wherein the transmit path channel multiplexer further comprises: a transmit channel equalizer.

16. The backhaul radio of claim 15, wherein the transmit channel equalizer produces transmit channel-equalized time domain samples by applying respective stream-specific or chain-specific transmit weights to the plurality of transmit symbol streams provided to the transmit channel equalizer.

17. The backhaul radio of claim 1, wherein a total number of receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays exceeds the number of the plurality of receive symbol streams.

18. The backhaul radio of claim 1, wherein a total number of receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays exceeds the number of said at least two of the plurality of receive RF chains that provide said at least two of the plurality of respective receive chain output signals to the frequency selective receive path channel multiplexer.

19. The backhaul radio of claim 18 further comprising one or more selectable RF connections for selectively coupling certain of the receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays to certain of the plurality of receive RF chains, wherein the number of receive antenna sub-arrays that can be selectively coupled to receive RF chains exceeds the number of receive RF chains that can accept receive RF signals from the one or more selectable RF connections.

20. The backhaul radio of claim 1, wherein each of the plurality of directional receive antenna arrays comprises two receive antenna sub-arrays.

21. The backhaul radio of claim 20, wherein a first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays has a first polarization, and wherein a second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays has a second polarization that is substantially orthogonal to the first polarization.

22. The backhaul radio of claim 21, wherein each of the plurality of receive antenna elements is one of a patch antenna element or a printed dipole antenna element.

23. The backhaul radio of claim 21, wherein said first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays comprises a plurality of patch antenna elements with each said patch antenna element excited in the first polarization, and wherein said second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays comprises the same said plurality of patch antenna elements with each said patch antenna element excited in the second polarization.

24. The backhaul radio of claim 23, wherein each of said plurality of patch antenna elements further comprises a first orthogonal probe feed for excitation in the first polarization for said first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays, a second orthogonal probe feed for excitation in the second polarization for said second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays, and a shorting pin at the center of each said patch antenna element.

25. The backhaul radio of claim 1, further comprising:

one or more probe-in-space antenna elements, wherein at least one of the one or more probe-in-space antenna elements has a wider antenna pattern beamwidth in azimuth than an antenna pattern beamwidth in azimuth of any of the receive antenna sub-arrays;

69

wherein at least one of the one or more probe-in-space antenna elements can be utilized for Dynamic Frequency Selection (DFS) that includes detection of radar systems.

26. The backhaul radio of claim 1, wherein each of the plurality of receive antenna sub-arrays is comprised of one or more columns of a plurality of patch antenna elements.

27. The backhaul radio of claim 26, wherein a number of columns for the one or more columns of a plurality of patch antenna elements is less than a number of patch elements within each of the one or more columns of a plurality of patch antenna elements.

28. The backhaul radio of claim 1, wherein each of the receive antenna sub-arrays comprised within each of the plurality of directional receive antenna arrays is oriented with respect to an azimuthal antenna pattern at a same azimuthal direction relative to at least one other of the receive antenna sub-arrays within its respective one of the plurality of directional receive antenna arrays and is oriented with respect to an azimuthal antenna pattern at a different azimuthal direction relative to at least one receive antenna sub-array within at least one other of the plurality of directional receive antenna arrays.

29. The backhaul radio of claim 1, wherein at least one of the plurality of directional receive antenna arrays is horizontally offset relative to the directional transmit antenna in a first direction and at least one other of the plurality of directional receive antenna arrays is horizontally offset relative to the directional transmit antenna in a second direction horizontally opposite to the first direction.

30. The backhaul radio of claim 29, wherein each of the receive antenna sub-arrays comprised within each of the plurality of directional receive antenna arrays horizontally offset relative to the directional transmit antenna in the first direction is oriented with respect to an azimuthal antenna pattern at an azimuthal direction that is substantially a mirror image about a center of an azimuthal antenna pattern of the directional transmit antenna relative to at least one receive antenna sub-array comprised within at least one of the plurality of directional receive antenna arrays horizontally offset relative to the directional transmit antenna in the second direction.

31. The backhaul radio of claim 1, wherein each of the plurality of directional receive antenna arrays is distributed in azimuthal orientation such that an aggregate normalized azimuthal antenna pattern beamwidth of the plurality of receive antenna sub-arrays in composite is equal to or greater than the antenna pattern beamwidth of the directional transmit antenna in azimuth.

32. The backhaul radio of claim 1, further comprising:

one or more modulator cores, wherein each one of the one or more modulator cores is capable of modulating one or more transmit data interface streams to produce one or more of a plurality of transmit symbol streams; and

a plurality of transmit RF chains, wherein each one of the plurality of transmit RF chains is capable of converting from one of a plurality of transmit chain input signals to a respective one of a plurality of transmit RF signals each comprising at least one of said plurality of transmit symbol streams;

wherein the directional transmit antenna is a directional transmit antenna array comprising a plurality of transmit antenna sub-arrays, and wherein each of said plurality of transmit antenna sub-arrays comprises a plurality of transmit antenna elements;

wherein each of said plurality of transmit antenna sub-arrays is coupled or coupleable to at least one transmit RF chain; and

70

wherein all of the plurality of directional receive antenna arrays are horizontally offset relative to said directional transmit antenna array.

33. The backhaul radio of claim 32, wherein the backhaul radio is capable of radiating at least one transmit RF signal from at least one of the plurality of transmit antenna sub-arrays at the same time that said at least two of the plurality of receive RF chains are capable of providing said at least two of the plurality of respective receive chain output signals to the frequency selective receive path channel multiplexer.

34. The backhaul radio of claim 32, wherein an antenna pattern beamwidth of each of the receive antenna sub-arrays is similar in elevation to an antenna pattern beamwidth of each of the transmit antenna sub-arrays in elevation.

35. The backhaul radio of claim 34, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays is narrower in azimuth than an antenna pattern beamwidth of each of the transmit antenna sub-arrays in azimuth.

36. The backhaul radio of claim 34, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays is narrower in elevation than an antenna pattern beamwidth of each of the receive antenna sub-arrays in azimuth.

37. The backhaul radio of claim 32, wherein each of the plurality of transmit antenna elements is one of a patch antenna element or a printed dipole antenna element.

38. The backhaul radio of claim 32, wherein each one of the one or more modulator cores comprises at least an encoder and a symbol mapper.

39. The backhaul radio of claim 32, wherein each one of the plurality of transmit RF chains comprises at least two digital to analog converters and a vector modulator that are capable of operating on the respective one of the plurality of transmit chain input signals, each said respective one of the plurality of transmit chain input signals comprised of digital baseband quadrature signals.

40. The backhaul radio of claim 32 further comprising:

a transmit path channel multiplexer, interposed between the one or more modulator cores and the plurality of transmit RF chains, to produce the plurality of transmit chain input signals provided to the plurality of transmit RF chains from the plurality of transmit symbol streams, wherein the transmit path channel multiplexer operates on one or more successive blocks of symbols from the plurality of transmit symbol streams.

41. The backhaul radio of claim 40, wherein the transmit path channel multiplexer comprises:

at least one cyclic prefix adder that prepends or appends a fraction of a block of transmit chain time domain samples to be transmitted sequentially in time to such block of transmit chain time domain samples to produce a cyclically-extended block of transmit chain time domain samples.

42. The backhaul radio of claim 40, wherein the transmit path channel multiplexer further comprises:

a transmit channel equalizer.

43. The backhaul radio of claim 42, wherein the transmit channel equalizer produces transmit channel-equalized time domain samples by applying respective stream-specific or chain-specific transmit weights to the plurality of transmit symbol streams provided to the transmit channel equalizer.

44. The backhaul radio of claim 32, wherein an antenna pattern beamwidth of each of the receive antenna sub-arrays is similar in elevation to an antenna pattern beamwidth of each of the transmit antenna sub-arrays in elevation.

45. The backhaul radio of claim 32, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays

71

is narrower in azimuth than an antenna pattern beamwidth of each of the transmit antenna sub-arrays in azimuth.

46. The backhaul radio of claim 32, wherein the antenna pattern beamwidth of each of the receive antenna sub-arrays is narrower in elevation than an antenna pattern beamwidth of each of the receive antenna sub-arrays in azimuth. 5

47. A backhaul radio comprising:

a vertically arranged backhaul radio antenna array comprising:

a plurality of directional receive antenna arrays, wherein each of said plurality of directional receive antenna arrays comprises a plurality of receive antenna sub-arrays, and wherein each of said plurality of receive antenna sub-arrays comprises a plurality of receive antenna elements; and 10

a directional transmit antenna; 15

one or more demodulator cores, wherein each demodulator core is capable of demodulating one or more of a plurality of receive symbol streams to produce one or more receive data interface streams; 20

a plurality of receive radio frequency (RF) chains, wherein each receive RF chain is capable of converting from one of a plurality of receive RF signals to a respective one of a plurality of receive chain output signals; and

a frequency selective receive path channel multiplexer, interposed between the one or more demodulator cores and at least two of the plurality of receive RF chains, wherein the frequency selective receive path channel multiplexer is capable of generating the plurality of receive symbol streams from at least two of the plurality of receive chain output signals provided by said at least two of the plurality of receive RF chains; 25 30

wherein each of the plurality of receive antenna sub-arrays is coupled or couplable to at least one receive RF chain;

wherein each of the receive antenna sub-arrays has an antenna pattern beamwidth in elevation that is narrower than an antenna pattern beamwidth in azimuth for any of the receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays; and 35

wherein all of the plurality of directional receive antenna arrays are vertically offset relative to the directional transmit antenna. 40

48. The backhaul radio of claim 47, wherein the directional transmit antenna has an antenna pattern beamwidth in elevation that is narrower than an antenna pattern beamwidth in azimuth for any of the receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays. 45

49. The backhaul radio of claim 47, wherein the backhaul radio is capable of radiating at least one transmit RF signal from the directional transmit antenna at the same time that said at least two of the plurality of receive RF chains are capable of providing said at least two of the plurality of respective receive chain output signals to the frequency selective receive path channel multiplexer. 50

50. The backhaul radio of claim 47, wherein at least one of the plurality of directional receive antenna arrays is vertically offset relative to the directional transmit antenna in a first direction and at least one other of the plurality of directional receive antenna arrays is vertically offset relative to the directional transmit antenna in a second direction vertically opposite to the first direction. 55 60

51. The backhaul radio of claim 47, further comprising:

at least one additional directional transmit antenna, wherein each one of the directional transmit antenna and the at least one additional directional transmit antenna is a directional transmit antenna array, wherein each of said directional transmit antenna arrays comprises a plu- 65

72

ality of transmit antenna sub-arrays, and wherein each of said plurality of transmit antenna sub-arrays comprises a plurality of transmit antenna elements;

one or more modulator cores, wherein each one of the one or more modulator cores is capable of modulating one or more transmit data interface streams to produce one or more of a plurality of transmit symbol streams; and

a plurality of transmit RF chains, wherein each one of the plurality of transmit RF chains is capable of converting from one of a plurality of transmit chain input signals to a respective one of a plurality of transmit RF signals each comprising at least one of said plurality of transmit symbol streams; and

wherein each of said plurality of transmit antenna sub-arrays is coupled or couplable to at least one transmit RF chain; and

wherein all of the plurality of directional receive antenna arrays are vertically offset relative to all of said directional transmit antenna arrays.

52. The backhaul radio of claim 51, wherein each of the receive antenna sub-arrays has an antenna pattern beamwidth similar in elevation to an antenna pattern beamwidth of each of the transmit antenna sub-arrays in elevation.

53. The backhaul radio of claim 51, wherein each of the plurality of directional receive antenna arrays is distributed in azimuthal orientation such that an aggregate normalized azimuthal antenna pattern beamwidth of the plurality of receive antenna sub-arrays in composite is equal to or greater than the antenna pattern beamwidth of each of the transmit antenna sub-arrays in azimuth.

54. The backhaul radio of claim 51 further comprising:

a transmit path channel multiplexer, interposed between the one or more modulator cores and the plurality of transmit RF chains, to produce the plurality of transmit chain input signals provided to the plurality of transmit RF chains from the plurality of transmit symbol streams, wherein the transmit path channel multiplexer operates on one or more successive blocks of symbols from the plurality of transmit symbol streams.

55. The backhaul radio of claim 54, wherein the transmit path channel multiplexer comprises:

at least one cyclic prefix adder that prepends or appends a fraction of a block of transmit chain time domain samples to be transmitted sequentially in time to such block of transmit chain time domain samples to produce a cyclically-extended block of transmit chain time domain samples.

56. The backhaul radio of claim 54, wherein the transmit path channel multiplexer further comprises:

a transmit channel equalizer.

57. The backhaul radio of claim 56, wherein the transmit channel equalizer produces transmit channel-equalized time domain samples by applying respective stream-specific or chain-specific transmit weights to the plurality of transmit symbol streams provided to the transmit channel equalizer.

58. The backhaul radio of claim 51, wherein each of the plurality of directional receive antenna arrays is distributed in azimuthal orientation such that an aggregate normalized azimuthal antenna pattern beamwidth of the plurality of receive antenna sub-arrays in composite is equal to or greater than an aggregate normalized azimuthal antenna pattern beamwidth of the plurality of transmit antenna sub-arrays in composite.

59. The backhaul radio of claim 47, wherein a total number of receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays exceeds the number of the plurality of receive symbol streams.

73

60. The backhaul radio of claim 47, wherein a total number of receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays exceeds the number of the plurality of receive RF chains that provide the frequency selective receive path channel multiplexer with the at least two of the plurality of respective receive chain output signals.

61. The backhaul radio of claim 60 further comprising one or more selectable RF connections that selectively couple certain of the receive antenna sub-arrays comprised within the plurality of directional receive antenna arrays to certain of the plurality of receive RF chains and wherein the number of receive antenna sub-arrays that can be selectively coupled to receive RF chains exceeds the number of receive RF chains that can accept receive RF signals from the one or more selectable RF connections.

62. The backhaul radio of claim 47, further comprising: one or more probe-in-space antenna elements, wherein at least one of the one or more probe-in-space antenna elements has a wider antenna pattern beamwidth in azimuth than an antenna pattern beamwidth in azimuth for any of the receive antenna sub-arrays;

wherein at least one of the one or more probe-in-space antenna elements can be utilized for Dynamic Frequency Selection (DFS) that includes detection of radar systems.

63. The backhaul radio of claim 47, wherein all of the plurality of directional receive antenna arrays are vertically offset relative to the directional transmit antenna in a same direction.

64. The backhaul radio of claim 47, wherein each of the receive antenna sub-arrays comprised within each of the plurality of directional receive antenna arrays is oriented with respect to an azimuthal antenna pattern at a same azimuthal direction relative to at least one other of the receive antenna sub-arrays within its respective one of the plurality of directional receive antenna arrays and is oriented with respect to an azimuthal antenna pattern at a different azimuthal direction

74

relative to at least one receive antenna sub-array within at least one other of the plurality of directional receive antenna arrays.

65. The backhaul radio of claim 47, wherein each of the plurality of directional receive antenna arrays comprises two receive antenna sub-arrays.

66. The backhaul radio of claim 65, wherein a first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays has a first polarization, and wherein a second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays has a second polarization that is substantially orthogonal to the first polarization.

67. The backhaul radio of claim 66, wherein each of the plurality of receive antenna elements is one of a patch antenna element or a printed dipole antenna element.

68. The backhaul radio of claim 66, wherein said first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays comprises a plurality of patch antenna elements with each said patch antenna element excited in the first polarization, and wherein said second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays comprises the same said plurality of patch antenna elements with each said patch antenna element excited in the second polarization.

69. The backhaul radio of claim 68, wherein each of said plurality of patch antenna elements further comprises a first orthogonal probe feed for excitation in the first polarization for said first one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays, a second orthogonal probe feed for excitation in the second polarization for said second one of the two receive antenna sub-arrays in each of the plurality of directional receive antenna arrays, and a shorting pin at the center of each said patch antenna element.

* * * * *